DOCUMENT RESUME

ED 344 605 IR 054 009

AUTHOR Michelson, Avra
TITLE Expert Systems Technology and Its Implication for Archives. National Archives Technical Information Paper No. 9.
INSTITUTION National Archives and Records Administration, Washington, DC.
PUB DATE Mar 91
NOTE 47p.
PUB TYPE Information Analyses (070) -- Viewpoints (Opinion/Position Papers, Essays, etc.) (120) -- Reports - Descriptive (141)

EDRS PRICE MF01/PC02 Plus Postage.
DESCRIPTORS *Archives; Business; Data Processing; *Expert Systems; Federal Government; Government Libraries; Information Storage; Information Systems; *Information Technology; Library Administration; Library Collections; National Libraries; Problem Solving; Public Agencies; *Records Management

ABSTRACT

This report introduces archivists to the potential of expert systems for improving archives administration and alerts them to ways in which they can expect intelligent technologies to impact federal record-keeping systems and scholarly research methods. The report introduces the topic by describing expert systems used in three Fortune 500 companies. It then defines expert systems, distinguishes them from conventional programs, and presents the capabilities of the technology together with examples of suitable applications. Discussion of the building of an expert system application begins with a short history of the evolution of the technology, followed by a detailed account of knowledge engineering, i.e., the process used to develop an expert system. Descriptions of several expert systems applications in the federal government highlight applications in the Internal Revenue Services, the Social Security Administration, and the Office of Management and Budget (Executive Office of the President). A report on the library profession's emerging use of this technology focuses largely on the three national libraries of the federal government: the National Library of Medicine, the National Agricultural Library, and the Library of Congress. The discussion of recent advances in expert systems technology that concludes the paper examines limitations of the technology, identifies likely frontiers for further research and development, and considers the implications of the technology for archives administration. A list of sources and related bibliographies is appended. (MAB)

****
* Reproductions supplied by EDRS are the best that can be made *
* from the original document. *
****

#2H18569

U.S. DEPARTMENT OF EDUCATION
Office of Educational Research and Improvement
EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)

☐ This document has been reproduced as received from the person or organization originating it.

☐ Minor changes have been made to improve reproduction quality.

• Points of view or opinions stated in this document do not necessarily represent official OERI position or policy

CD
9T1
.N35
No.9

ED344605

EXPERT SYSTEMS TECHNOLOGY

AND

ITS IMPLICATION FOR ARCHIVES

ARCHIVES
LIBRARY
INFORMATION
CENTER

National Archives Technical Information Paper No. 9
March 1991

Avra Michelson
Archival Research and Evaluation Staff
National Archives and Records Administration

"PERMISSION TO REPRODUCE THIS MATERIAL HAS BEEN GRANTED BY

Jeffery T. Hartley

TO THE EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)."

BEST COPY AVAILABLE

1205400

TABLE OF CONTENTS

PREFACE iii

ACKNOWLEDGEMENTS iv

I. Introduction 1

II. Building Expert Systems : History 6

III. Building Expert Systems: Knowledge Engineering 7

IV. Federal Government Use of Expert Systems 13

V. The Library Profession's Use of Expert Systems 18

VI. State of the Technology 22

VII. Implications for Archives 27

SOURCES 32

BIBLIOGRAPHY 35

LIST OF FIGURES

Expert System Architecture 9

Examples of Rules 11

National Archives

Washington, DC 20408

March 5, 1991

This report, *Expert Systems Technology and Its Implications for Archives*, explores the potential application of expert systems technology to archival programs and processes. It both introduces archivists to a new software tool that may be used to improve archival operations, and alerts them to a technology successfully used in business, industry, government agencies, and increasingly by scholars.

The report represents the culmination of eighteen months of research on artificial intelligence and expert systems. In preparing the report, members of the Archival Research and Evaluation Staff of the National Archives received briefings from professional staff of many of the country's most advanced artificial intelligence research laboratories. They also met with managers of operational expert systems applications, observed demonstrations of applications in use or underway, evaluated relevant literature, and participated in educational training sponsored by professional associations. This final version of the report was reviewed by a team of technical experts.

By releasing this report to the public, the National Archives invites an exchange of views with the archival community and related professions on the implications of the technology. We look forward to your comments on the report and the opportunity to explore the potential of expert systems technology in archival administration.

DON W. WILSON
Archivist of the United States

PREFACE

The National Archives and Records Administration (NARA) acquires, preserves, and makes available for use the permanently valuable records of the United States Government. Within NARA, the Archival Research and Evaluation Staff evaluates the ways emerging technologies can be used to improve the management of archival records. The staff assesses technologies through the development of technical information reports and pilot/prototype projects, or through sample evaluations of new products or systems.

This report introduces archivists to the potential of expert systems for improving archives administration. Further, it alerts archivists to ways in which they can expect intelligent technologies to impact federal record-keeping systems and scholarly research methods.

The report introduces the topic by describing expert systems use in three Fortune 500 companies. Next, it defines expert systems and distinguishes them from conventional programs. The capabilities of the technology are presented, with examples of suitable applications.

Section two and three considers the building of an expert system application. It opens with a short history on the evolution of the technology, followed by a detailed account of knowledge engineering - the process used to develop an expert system.

The fourth section reviews several expert systems applications in the federal government. The discussion of federal government applications features the Internal Revenue Service, the Social Security Administration, and the Office of Management and Budget (Executive Office of the President). A report on the library profession's emerging use of this technology appears next, focusing largely upon the three national libraries of the federal government - the National Library of Medicine, the National Agricultural Library, and the Library of Congress.

A discussion of recent advances in expert systems technology is the focus of the final portion of the technical information paper. This section examines limitations of the technology and identifies likely frontiers for further research and development. The technical information paper concludes by considering the implications of the technology for archives administration.

ACKNOWLEDGEMENTS

Many individuals, both in government service and the private sector, contributed to the development of this report. Unfortunately, we are unable to formally recognize the dozens of people who generously provided technical advice, research leads, and information on applications. We appreciate their assistance and hope that we may return the favor some day.

The National Archives would like to extend special thanks to those who spoke with us about their expert systems applications, to the artificial intelligence specialists consulted in the process of research, and especially to Dr. Rick Steinhauser for his kind invitation to participate in a government tour of artificial intelligence research laboratories, and to Jeff Rothenberg, Harry Siegel, and Ted Senator for their valuable comments as technical readers of the report.

EXPERT SYSTEMS TECHNOLOGY AND ITS IMPLICATION FOR ARCHIVES

I. Introduction

In 1956, the Institute of Radio Engineers sponsored a symposium on the impact of computers on science and society.[1] The symposium brought together many pioneers of computer technology to explore the technical, social, and political implications of their work. One speaker, a researcher from International Business Machines, Inc. (IBM), expressed optimism about the "first groping steps" the corporation was taking toward endowing machines with some sort of intelligence. He went on to comment on what he considered the real promise of computer technology:

> "The real importance of machines lies in their capacity as general logical devices. It's very easily understandable why, during these first ten years of computer technology, we have thought of them generally as calculating devices, but it is equally clear...that the really big problems, the problems which can greatly affect scheduling of production, with the planning and control of traffic in an airline,... or in general with the efficient conduct of any complex activity, will require on our part a rather different technique of machine use than we have yet developed....So far, we have all...tended to think of computers as essentially mathematical computing devices....It's a rather young attitude on our part....In my opinion, this is only the beginning."[2]

The observations of this young researcher turned out to be prophetic. Although it took thirty years, corporations and federal agencies now use "smart" technologies to assist with activities that involve human reasoning, such as the scheduling, planning, and managing of complex tasks. During the past decade, the trend in computer technology toward smaller hardware, improved memory capacity, and decreased costs has encouraged the widespread automation of workplaces, and even private homes. As a result, to a great extent computers already handle much of the travail of the information age: the sorting and storing of data. To a surprising degree, computers now perform jobs once considered impenetrable human preserves: executing tasks and solving problems that involve logical reasoning using stored information.

Expert systems are computer applications that can consider a vast amount of knowledge, then "reason" and recommend a course of action. No longer considered an emerging technology, expert systems have advanced from the research bench to operational maturity. During the past decade, most Fortune 500 companies introduced expert systems technology into corporate operations in an attempt to secure a competitive edge. IBM began commercializing the technology in 1985 and now has nearly 400 major expert systems

---

[1] "Symposium on the Impact of Computers on Science and Society," held at the 1956 National Convention of the Institute of Radio Engineers, Philadelphia (Waldorf-Astoria), March 22, 1956.

[2] Remarks of Dr. David Sayre, "Symposium on the impact of computers on science and society," *IRE Transactions on Electronic Computers*, Volume EC-5/3 (September 1956), p. 149.

applications in use or under development, with six of the leading applications resulting in a net return in 1989 of $38.5 million. For example, one of the applications, CASES, functions as a support tool for staff who acquire and move equipment. It performs this task by identifying the particular forms that require completion and the signatures needed, and by offering precautionary advice and related information to the user. Since its introduction into operations in 1986, CASES has significantly reduced the number of routine questions posed to finance managers, in addition to reducing signature authorization problems and noncompliance with procedures. The system was developed in a year and a half at a cost of less than $100,000.

E.I. DuPont integrated expert systems technology into its operations by internally developing many small applications. By 1989, the corporation reported 600 expert systems in use with an average annual return of 10 to 1 on each system. Since 1985, DuPont has used expert systems to assist in three types of tasks: 1) diagnosis - for example, determining why a product sold to a customer doesn't work, why a manufacturing yield is poor, or which part of a piece of equipment is responsible for a malfunction; 2) selection - identifying which of the 600 widgets manufactured by the company is most appropriate for a particular customer, or the correct course of action when one of 200 chemicals leaks into the environment; and 3) planning/scheduling - reconciling contention for machinery while minimizing loss and production lapses within specified time constraints. Despite the high monetary ratio of return, DuPont reports that the greatest payback on its investment results from the replication of expertise - the freeing of experts on staff to perform more difficult jobs, by transferring task responsibility from experts to non-experts.

American Express Company (AE) offers a third example of corporate use of expert systems technology. They built a system that assists staff in authorizing the credit-worthiness of customers with marginal credit histories. In the actual work situation, AE staff are required to make each credit determination in less than 70 seconds. The expert system AE developed assists staff in this task by analyzing a customer's financial record to arrive at a credit approval recommendation for a particular purchase. The company estimates that the $4 million investment in the system protects it from an annual $27 million loss from bad credit and fraud and has reduced erroneous judgments on purchase approvals by about 75 percent. Corporate managers report, however, that the system's greatest contribution is that it simultaneously performs the normally contradictory activities of improving customer service while reducing losses. Optimizing sensible credit approvals in marginal cases allows the company to maximize its profits.

These three companies use expert systems, one of several intelligent technologies,[3] to perform tasks normally thought to fall outside the range of machine capabilities. Business and industry are not the only enterprises to discover these technologies. Military and intelligence agencies, as well as countless civilian government offices, use computers to solve real-world problems that involve human judgments. In fact, forecasters anticipate that the near-term social impact of intelligent technologies will compare to that of the industrial revolution. Over the course of a century, the industrial revolution irrevocably transformed

---

[3] Subfields of artificial intelligence are considered "intelligent" technologies. See footnote four for further explanation.

the entire spectrum of physical labor. Today, intelligent machines are revolutionizing the work of the mentally-skilled: professional and semi-professional staff who perform with their minds. Expert systems, software technology that assists the human mind, are being used to enhance the productivity and skill of intellectual labor in much the same way that great machines first assisted human muscles in the last century.

What are expert systems?

Expert systems are computer tools that help humans complete a task, solve a problem, or make a decision that requires logical reasoning. As a subfield of artificial intelligence (AI),[4] the key characteristic of expert systems lies in its use of judgments (or human-like reasoning) instead of mathematical calculations (or algorithms) to execute procedures. This feature distinguishes expert systems from conventional programming. They formulate solutions by using factual knowledge as well as heuristics, "rules of thumb" that establish the basis for an educated guess. Unlike algorithms, a heuristic approach to problem solving entails judgments and often a degree of uncertainty. The ability to make machines perform with human-like reasoning and uncertain logic represents the chief contribution of expert systems technology to computer science.

Expert systems function as assistants rather than replacements. Most are designed to handle about 85 percent of the cases for a particular domain. They can be developed as stand-alone applications, as embedded code in conventional applications, or as systems fully integrated in operational environments.

In a typical application, the system first requests initial information about a problem. It then searches the knowledge stored in the computer for a rule, pattern, or model that "matches" characteristics of the problem suggested by the initial data. The application may offer an immediate solution or request more information. This process continues until the system arrives at a conclusion or recommends an action. Ideally, such systems can explain to users at any time the line of questioning, the reasoning behind a particular conclusion, or the relevance of a suggested course of action.

The effectiveness of expert systems performance has been established with a variety of intellectual tasks. Some of the most productive systems have been developed for tasks such as 1) scheduling/planning - applications that integrate actions, events, and objects within specific time and resource constraints; 2) diagnosis - applications that identify diseases, organic substances, machinery problems, and so forth, based on known information;

---

[4] Artificial intelligence is a subfield of computer science concerned with developing machines capable of performing functions normally associated with human intelligence, such as reasoning, learning, and understanding human language. Some of the major application areas within the field are robotics, machine vision, machine voice recognition, natural language processing, expert systems, and neural networks. The pursuit of machines capable of performing tasks that require "intelligence" has been a source of controversy. Opponents argue that human intelligence can never be reduced to machine processing, hence the goal of producing "intelligent" machines is not realizable. This paper does not engage the question of whether or not AI applications possess reasoning capabilities equivalent to human intelligence because the issue is irrelevant for federal program managers. As government staff, the only important issue in considering AI technology is whether or not it works -- does it solve an operational problem?

3) computer-aided learning - applications that use computers to transfer expertise from senior to junior staff for training purposes; 4) selection - applications that classify objects and recommend a course of action appropriate to the classification; 5) data interpretation - applications that infer patterns and anomalies, generate likely outcomes, and discover new information with large, complex aggregates of data; and 6) user interfaces - applications that support query negotiation and information retrieval in an automated environment.

An impressive range of applications already has been developed. The financial community was one of the earliest professions to adopt expert systems, implementing the technology with portfolio management, with risk analysis, and to evaluate major strategic investments. Manufacturers use the technology to control production processes, diagnose system malfunctions, and design complex machinery. Banks and insurance companies use expert systems for asset management, credit analysis, medical assessments, and overall management functions. Expert systems also provide on-site managerial assistance to retail chain managers.

The scientific professions have developed expert systems for many purposes, including diagnosing diseases, providing advice on scientific experiments, and helping unskilled metallurgy staff identify commercially used compounds. In the transportation industry, expert systems control the handling, labeling, and shipping of hazardous chemicals. Automobile manufacturers, such as Toyota, use the technology to diagnose engine problems in passenger cars and to customize vehicle design. Expert systems assist cartographers in compiling features on nautical maps and charts.

Examples of expert systems are almost as numerous as varieties of professional and semi-professional work. Expert systems are used extensively to provide legal advice, to support geological analysis, for marketing purposes, for the analysis of actuarial data, to make social service eligibility determinations, for military battle management, for evaluation of loan applications, and even by the Chinese to capture the ancient knowledge of rural herbal doctors.

Although at last count more than 2200 operational applications existed in the business world, and the estimates for the number of government applications are even greater,[5] archivists are just beginning to examine the feasibility of using the technology to solve mission problems. In 1985, the National Archives and Records Administration (NARA) developed a proof-of-concept prototype expert system that emulates the reasoning strategies of reference

---

[5] Reliable figures on the extent of operationalized expert systems are unavailable. *Expert System Strategies* (5:1 Jan. 1989) estimated that in 1989 there were 2200 business applications in the U.S. Editor Paul Harmon predicted, based on the number of prototypes underway and the number of new companies acquiring the technology, that this number would double in 1990. Feigenbaum, et al., reported 1500 business expert systems in the U.S. in 1988. Neither of these estimates includes government expert systems, a significant portion of the market. Indeed, the Commerce Department's *Industrial Outlook 1989* reports that the federal government accounts for 75 percent of the U.S. expert systems market. Although the federal government does not track use of the technology, it is commonly recognized that most applications have been developed by intelligence agencies, the military, or the space program. However, the number of civilian government applications is rapidly increasing.

archivists in negotiating patron requests. Although the prototype successfully demonstrated proof of concept, NARA did not pursue further research or development at that time.[6]

The only other archival application of the technology is found in a system that helps archivists identify objectives for strategic planning in preservation management. The National Association of Government Archives and Records Administrators (NAGARA) sponsored the development of a stand-alone application, which includes approximately 320 production rules developed using EXSYS software. The system is currently out for pilot testing.[7] A related research project that will be carried out during the next four years in conjunction with the Swiss Federal Archives investigates the use of natural language processing to support automatic filing of documents in an electronic environment.[8] Natural language processing, another subfield of artificial intelligence, involves the development of machines capable of understanding language as people typically speak and write it.

New and emerging technologies promise to modernize the performance of archival administration tasks and impact the conduct of standard archival practice. This report introduces archivists to the ways intelligent technologies, and in particular expert systems, can be used to improve the management of archives. The studies conducted with operational expert systems reported in this paper clearly demonstrate that the technology improves speed, productivity, and consistency when used with many professional and semi-professional tasks. Further, widespread use of the technology by federal agencies, and future use by scholarly researchers, may make current aspects of archival practice obsolete. For example, most information systems developed by the federal government do not include a complete record of the rules and regulations that comprise an agency's policy. In the future, however, an agency's policy likely will appear increasingly as code in intelligent programs. This trend suggests the need to reassess conventional appraisal practices. Moreover, forecasters predict that once widely available, scholars will adopt "smart" software to inform their analyses. For these reasons archivists need to understand intelligent technologies, begin to integrate them into appropriate operations, and consider the operational implications of their use by federal agencies and scholars.

---

6 The goal of the project was to apply the findings of the reference archivist proof-of-concept expert system to the design of an automated information system under development for the Office of the National Archives. According to NARA's Five-Year ADP Plan, an expert system front-end interface for this system is scheduled for development around 1994. For further information on the proof-of-concept system, see: Renee M. Jaussaud, "The archivist's assistant: from the expert's perspective," in *Managing Artificial Intelligence and Expert Systems*; Daniel A. De Salvo et al., "Structured Design of an Expert System Prototype at the National Archives," in *Expert Systems for Business*; and Daniel A. De Salvo and Jay Liebowitz, "The application of an expert system for information retrieval at the National Archives," *Telematics and Informatics*. (For the full citations, see section 6 of the bibliography).

7 The NAGARA expert system was developed by Bonnie Curtin, Georgia Department of Archives and History, as one part of a three-part archival preservation planning instrument.

8 Funding for the project is provided by *Fonds national suisse de la recherche scientifique*. Christoph Graf is the project director at the Swiss Federal Archives. The principal AI investigator is Margaret King.

II. Building Expert Systems : History

For more than 30 years, artificial intelligence researchers have explored human cognition in an effort to develop machines capable of emulating human thinking. The earliest attempt to develop an intelligent program resulted in the General Problem Solver (GPS). Investigators built the GPS by extracting general strategies for problem solving from interviews in which people were asked to solve problems out loud. This research typified the AI community's emphasis during the first decade of research on discovering generic human reasoning strategies. After many years of trial applications, a new tenet, known as the knowledge principle, emerged as the fundamental principle for creating machine "intelligence." The knowledge principle suggests that the quality and quantity of knowledge coded in a program is the key to achieving high levels of machine competence with complicated intellectual problems. Reasoning strategies, although necessary, have secondary importance for machine problem solving. Expert systems emerged in the 1960's and 1970's as the practical application of the knowledge principle. By the end of the 1980's, expert systems technology achieved operational maturity.

One of the earliest successful systems, Dendral, assists chemists to determine molecular structures by analyzing spectrograph data. Edward Feigenbaum and his colleagues at Stanford University spent more than 15 work-years developing Dendral. Today it is found in organic chemistry laboratories throughout the world, and its use has resulted in more than 50 publications in the chemistry literature.

During this same period, Massachusetts Institute of Technology (MIT) researchers created Macsyma, an application that tackles mathematical problems including calculus, algebra, and the solving of equations. After nearly 100 work-years of development, Macsyma performs better than most human experts in the symbolic conduct of differential and integral calculus.

Mycin, developed at Stanford in 1972, diagnoses bacterial blood infections and recommends medical treatment. Emycin, an expert system software development tool, evolved from Mycin around 1978. It was named Emycin (Empty Mycin) because it contains the logical structure of Mycin without its specific knowledge of blood disorders. Emycin marks the development of the first expert system shell - a program with logical structures and inferencing strategies, but without the specific knowledge required for a particular application.

Also during the 1970's, the Stanford Research Institute developed Prospector, an application designed to advise on mineral prospects of a region by producing favorability maps. Funded by the U.S. Geological Survey (USGS) and the National Science Foundation, Prospector provided consultation to geologists in the early stages of identifying probable ore-grade-deposit sites. Since 1983, the USGS's successor system, Prospector II, has evolved into a full-sized system to aid geologists to identify regional mineral deposit types.[9]

---

[9] For further information on Prospector II see: Richard B. McCammon, "Maintaining Prospector II as a full-sized knowledge based system." *AISIG '90 Research Workshop: Working Papers on Full-Size Knowledge Based Systems.* (May 7-8, 1990), pp. 1/14 -1/32.

The first commercial operational expert system, Xcon, was developed by Digital Equipment Corporation (DEC) in 1980. Xcon configures the company's response to customers' requests for VAX computer systems by reconciling space requirements, client specifications, and technical constraints. DEC estimates that the application saves the corporation approximately $200,000 per month in staff costs, as well as an uncalculated amount in manufacturing costs.

These early expert systems constitute only a fraction of the thirty-year research effort expended to develop an operational technology. Although the systems span a variety of domains, a common lesson emerged from the development efforts. That lesson - the paramount importance of knowledge - formed the basis for all subsequent achievements in today's expert systems technology.

III. Building Expert Systems: Knowledge Engineering

Expert systems are designed and built by knowledge engineers, computer scientists trained in applied artificial intelligence methods. Knowledge engineers develop expert systems in close coordination with domain experts, individuals who by training, education, or experience have demonstrated proficiency within a particular field. The chief quality that distinguishes domain experts from others in a craft is expertise - the capacity to efficiently and effectively solve discipline-specific problems.

Knowledge engineers use the expertise acquired from experts to develop systems in a particular domain. Domains are fields of knowledge in which humans use intelligent judgments to reach a decision or complete a task. Archives administration could be considered a broad domain, while records appraisal represents a narrower domain. Today's expert systems perform most effectively with narrow, well-bounded problems that involve judgments but not a lot of what is considered "common sense" to solve.

Expert systems typically consist of a knowledge base, an inference engine, an explanation facility, and a user interface. The knowledge base acts as the system's repository for facts and heuristic rules used by the expert to solve domain problems. Factual rules represent laws, regulations, and knowledge commonly agreed upon by experienced members of a field. In contrast, heuristic rules constitute judgmental, experimental, and often uncertain personal knowledge comprised of rules of thumb that form the basis for educated guesses. These procedural tips and "tricks of the trade" help to guide, limit, and speed up the search process. Most individuals invoke widely known heuristic rules when selecting wine to complement a meal; e.g., if the meal is chicken or fish, then select a white wine. This type of heuristic knowledge narrows the search, resulting in more efficient decision-making.

Unlike conventional programs, expert systems are characterized by an absolute distinction between knowledge and processing controls. Since the rules and facts found in expert systems reside in static knowledge bases, another mechanism must control where to start the reasoning process and which rule to examine next. For this reason, expert systems contain inference engines, the drivers of the rules that perform the reasoning function. The inference engine controls the process of invoking the rules necessary for solving a particular problem. Inference engines perform this function through a search strategy.

Many different search strategies are used with expert systems, but most are based on one of two fundamental concepts: forward chaining or backward chaining. Forward chaining is a bottom-up reasoning process that starts with known conditions and works towards a desired goal. Backward chaining is a top-down reasoning process that starts from a desired goal and works back to determine if the requisite supporting conditions exist. A machine asked to prove a theorem would use backward chaining, while determining the next move in a chess game involves a forward-chaining search.

Explanation facilities, another critical feature of expert systems, are not found in conventional computer programs. Unlike expert systems, conventional programs typically process data by repeatedly running the same predefined procedures that provide 100 percent certain results. For example, automated banking systems arrive at a daily balance by processing each day's deposits and withdrawals through the same set of computer procedures. Even though the set of numbers changes daily, the same computer procedures are used to arrive at a daily balance, eliminating the need for an explanation facility.

Expert systems, however, reach their recommendations by using uncertain and probabilistic logic with a unique series of rules. Moreover, the system's recommendations represent satisfactory solutions, not absolute answers as with conventional applications. Assuming more than one acceptable answer is possible, expert systems are programmed to justify their recommendations by supplying the rationale. It is this explanation facility that describes how decisions are reached and the steps taken to reach them. Its purpose is to increase user confidence in the system by supplying the information needed to understand and assess the system's logic.

The final standard component of expert systems is the user interface. As is customary with conventional programs, the interface facilitates the interaction of the computer program with the end user. Interfaces typically translate user communications into a form that can be processed by the rest of the system. They also convert the machine's processing into a form a user can understand. The interface for an expert system is comparable to those found on word processing software or bibliographic systems - it provides the user with a means to carry on a dialog with the system (see Figure 1).

Expert systems software can be developed either with special artificial intelligence hardware and programming languages, or on standard computers with either a commercial expert system software shell or in a standard software programming language such as "C". Computer scientists have developed two AI programming languages: LISP (List Processing), the language of choice in the United States, and PROLOG (Programming Logic), most popular in Japan and Europe. A LISP machine, special symbolic processing hardware, often is used with LISP to optimize programming performance. Although the programming environment with the LISP machine includes many desirable features currently unavailable with standard computers, the clear trend is toward the diminishing use of symbolic processors in favor of standard hardware.[10]

---

[10] Symbolic processors manipulate symbols that represent primarily logical rather than numerical operators. Today's manufacturers of symbolic processors are Symbolics, Inc. and Texas Instruments, Inc.

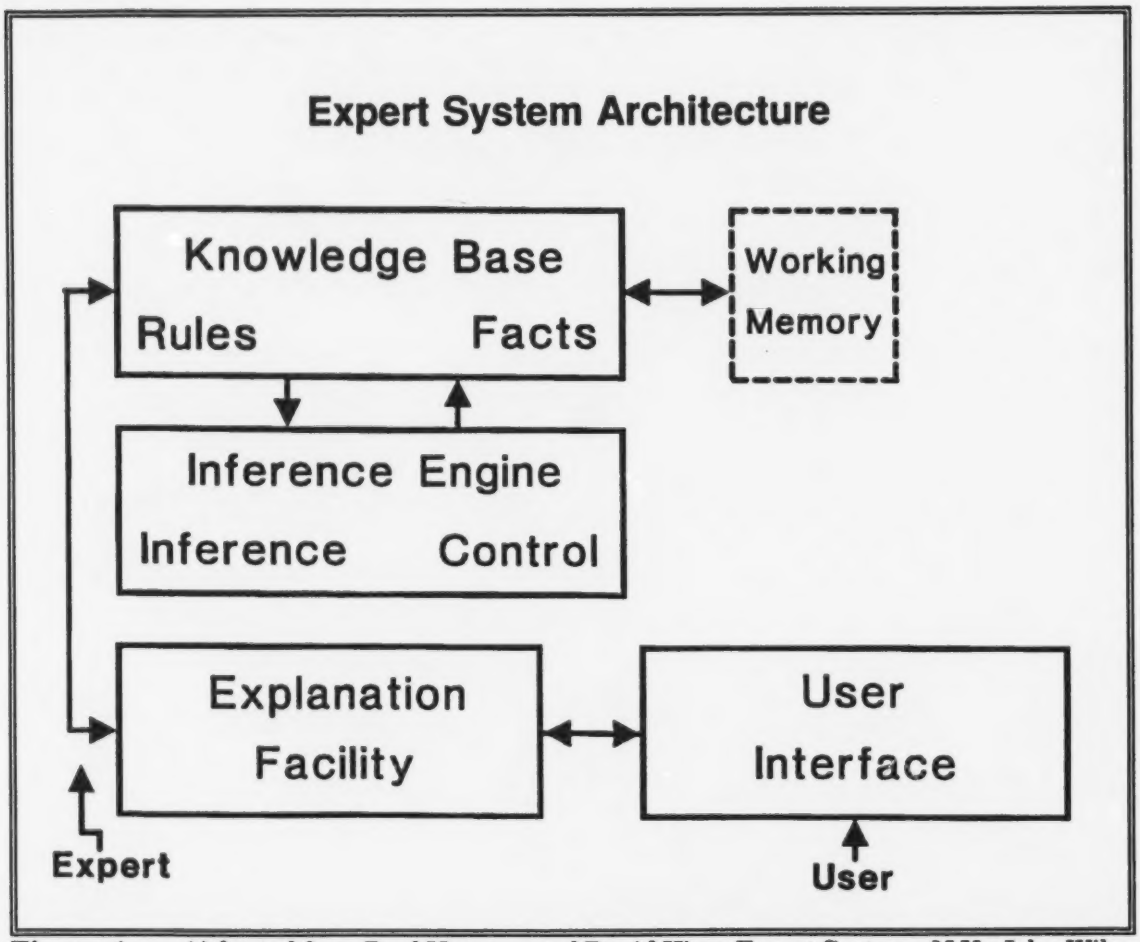

Figure 1 (Adapted from Paul Harmon and David King, *Expert Systems*, N.Y.: John Wiley & Sons, Inc., 1985, p. 34)

In today's market, most expert systems are developed with some sort of commercial shell. Expert system software shells are sometimes compared to prefabricated houses; they are a set of tools and templates for building an expert system that allow for some customization of the product. Off-the-shelf software tools provide skeleton expert systems that typically include a predefined inference engine with empty buckets for developing a customized knowledge base. They may also contain a validator program that alerts the knowledge engineer to contradictory facts and rules, or redundant and missing information, mechanisms for editing the knowledge base, correcting syntactical errors, or debugging the system during development, as well as spelling checkers, windowing packages, and integration procedures for use with conventional software.

Software shells characteristically reduce development time, but they necessarily limit development options. Typically, they support only one or two methods of inferencing and knowledge representation. When using software shells, knowledge engineers must accept either the limitations of a particular shell or expend considerable effort to program around

them. For this reason, many knowledge engineers prefer to build a system from scratch on symbolic processors with an AI language, then convert it to conventional hardware and software for use. This is particularly true for large, complex systems or high-level AI research applications. Software shells tend to be used more frequently by in-house developers and commercial vendors to build relatively straight-forward small to mid-range applications.

Whether starting from scratch or using a shell, knowledge engineers translate the expertise of experts into a language that a machine can manipulate. System development begins with the knowledge engineer eliciting from an expert or group of experts, the knowledge used to solve problems in a domain. The process of extracting, structuring, and organizing knowledge acquired from human experts is one of the most difficult, yet critical, aspects of expert systems development. Because expert systems emulate human intelligence, the quality of the knowledge reported by domain experts directly affects the caliber of the system's performance. Artificial intelligence researchers draw on cognitive psychology to explain the variety of problems encountered in eliciting domain knowledge from experts.

It is widely recognized, for instance, that experts often have difficulty communicating pertinent knowledge, especially 1) knowledge associated with complex tasks, 2) the full range of knowledge used for decision-making in a field, and 3) implicit knowledge - knowledge used repeatedly in a procedure until the procedure is performed automatically and the expert has lost the ability to remember and verbalize the underlying process (such as much of the knowledge used when driving a car).

Structured interviews and protocol analysis are two of the most common methods used by knowledge engineers to acquire knowledge from domain experts. The typical protocol procedure calls for domain experts to "think aloud" while making domain decisions. Using methods borrowed from cognitive psychology, a knowledge engineer learns what experts know and how they reason with their knowledge.[11]

Knowledge engineers translate the knowledge acquired from experts into a form usable by a computer. This transformation of human knowledge into computer-readable symbols is referred to as knowledge representation. Production rules are the most conventional form of symbolic knowledge representation. They consist of "if" and "then" parts; "if" lists a set of conditions presented in some logical combination, while "then" represents the conclusion of the "if". Expert systems that rely on production rules as the basis of representation are referred to as rule-based systems (see Figure 2).

---

[11] Renee M. Jaussaud's article cited in footnote 6 relates a knowledge acquisition experience from the perspective of the domain expert.

Examples of Rules

From NAGARA's strategic planning for preservation application:

IF: (15.1.8) Do you have a procedure for identifying and selecting archives records that require a security copy and/or use copy? NO

THEN: GOAL I, Reprography Step One, OBJECTIVE D (a) - Develop appraisal strategies that relate record reproduction options to use/value and original record media. (2.1.1.7.1) - Probability=9/10

From NARA 's Archivist's Assistant:

if geo = GEO
and geo_specific (GEO,1) = GEOSPECIFIC
and do (reset series (GEOSPECIFIC, RG)
then geo_search (RG, TOPIC)

A paraphrase of a rule from NARA's Archivist's Assistant:

IF the information-sought is suits_against_the_Dept._of_Interior
THEN examine record group 46 cf 70

Figure 2

Another common form of knowledge representation, frames, packages knowledge in a way that accommodates cognitive associations. Frame-based systems are noted for their efficient representation of detailed object relationships, such as the association that a living room typically contains a couch and chairs, but not an ocean. New methods of knowledge representation with alternate advantages, such as case-based reasoning and reasoning by analogy are beginning to migrate from AI laboratories to the knowledge engineer's workbench.

In addition to acquiring and representing knowledge, the knowledge engineer designs the reasoning structure that controls the search and execution of rules (or optimizes the inference engine used in a software shell). To complete the process, the knowledge engineer determines how to integrate the use of uncertain knowledge into the reasoning process and the types of explanations about a system's recommendations that would benefit the end user.

One of the key principles of knowledge engineering is iterative development practiced in a manner that incorporates extensive involvement of the expert(s) into all stages of building the program. Knowledge engineers have adopted an iterative approach to expert system design to overcome the difficulties associated with the cornerstone development activity:

knowledge acquisition. Knowledge acquisition is fundamentally a process of communicating the mental reasoning and thought processes of an expert to the knowledge engineer. But the problems associated with both the recollection and communication of complex and implicit knowledge are extensive.

After considerable research, knowledge engineers discovered that iterative prototype development can be used to substantially reduce the difficulties associated with knowledge acquisition. Iterative prototyping involves the successive repetition of knowledge acquisition and knowledge representation until a system reaches the desired goal. Systems unfold in stages, with each stage an improvement on the proficiency of the previous stage. This iterative method of development expedites knowledge acquisition because gaps in information and lapses in logic imperceptible to humans ultimately are discovered while building a workable application. For example, a domain expert may communicate to a knowledge engineer five steps for solving a problem. Because the expert's reasoning is so automatic, he/she inadvertently fails to express an important consideration. The domain expert may assure the knowledge engineer in good faith, however, that the procedures conveyed are complete. It is not until the knowledge engineer's program is run and found to fail that either member of the team is aware of the omission. Further, iterative prototyping points out the knowledge engineer's misunderstandings, misinterpretations, or misrepresentations of the expert's meanings. In short, knowledge engineers have discovered that in most instances the best technique for recognizing their own mistakes or for prompting experts' recall is to show experts their methods in operation. Consequently, expert systems are built through an iterative process of development.

It is difficult to suggest approximate costs for expert systems, as each application entails domain-specific problems that affect development times, hardware and software selection, and requisite staff resources. Although an inadequate gauge, expert systems usually are measured through the number of rules or rule-equivalents.[12] Small systems average from 50-500 rules and are typically developed on a microcomputer with a software shell. Mid-range systems generally contain 500-1500 rules and are developed on supermicrocomputers or minicomputers using a shell or a high-level programming language. Large systems exceed 1500 rules and are programmed for a minicomputer or mainframe environment with a high-level or AI language. Software tools appropriate for small applications average from about $100-$1,000. The mid-range application tools cost approximately $2,500-$25,000, with most selling between $5,000-$8,000. Mainframe software varies in price, but it averages about $60,000 - $100,000. As with conventional software, the price/capability ratio of expert systems software tools is decreasing.

---

[12] The "size" of an expert system is difficult to determine. Ted Senator, U.S. Navy, in a presentation to the IEEE Technical Committee on Expert Systems (WDC, June 6, 1990) suggests that the number of rules is an inadequate measure, because "large" systems simply could be indicative of sloppy knowledge engineering and responsible for a system's inefficiency. Further, as new forms of knowledge representation become more common, fewer expert systems contain rules per se. Senator proposed that the size of an expert system be evaluated by assessing the problem scope (the number of domains and sources of knowledge), the size of the knowledge base (including number of rules or rule equivalents, search space, and lines of code), the code required to embed or integrate the system into other operations, the performance expectations, the number of users and sites, and the ultimate deployment of the system.

IV. Federal Government Use of Expert Systems

Even though more than 80 percent of Fortune 500 companies are integrating expert systems into their operations, the government is still considered the greatest user of the technology. Indeed, the Commerce Department reports that approximately 75 percent of the U.S. expert systems market belongs to government.[13] In particular, military, intelligence, and the space program have sponsored artificial intelligence research and development in this country for decades. An Office of Technology Assessment projection estimated that in 1988 expenditures for defense-related artificial intelligence research would reach $155.5 million, while the government would devote only a fraction of that figure, $17 million, to civilian AI research.[14] For many years, the Defense Advanced Research Projects Agency (DARPA), the chief source of military funding for AI research, has funded research projects to develop military systems that assist in strategic planning and in determining a course of military action.

Today, the technological breakthroughs achieved with defense-oriented systems are finding application in civilian systems. There are no accurate figures, however, on how many civilian government expert systems applications exist. Neither the Office of Management and Budget nor the Congressional Budget Office tracks this type of expenditure across agencies. Moreover, it is difficult to establish the pattern of use within any one agency because the technology tends to be administered like any other mature software tool, and therefore, does not appear as a separate budget item. For these reasons, knowledgeable sources within the federal government hesitate to estimate formally the extent of civilian use; however, off the record they speculate that the development of prototypes and operational expert systems is currently "widespread and growing," and in the next 5-10 years, expect that "inevitably" they will be used throughout government.

The purpose of this section of the paper is to illustrate, through case examples, the nature and extent of use of expert systems technology within federal civilian agencies. This section provides archivists with a glimpse of the widespread use of the technology for a broad range of tasks. The status of the technology is examined in three agencies: the Internal Revenue Service, the Social Security Administration, and the Office of Management and Budget (Executive Office of the President). Applications underway in a half-dozen other agencies are also briefly mentioned. The sketches that follow introduce readers to a small portion of federal civilian applications already operational or under development.

Internal Revenue Service: Virtually every citizen, as a taxpayer, is subject to oversight by the Internal Revenue Service (IRS). The IRS is responsible for enforcing internal revenue laws and collecting the proper amount of taxes. It must perform these activities at the least cost to the public and in a manner that fosters the highest degree of confidence in the agency's integrity, efficiency, and fairness. During the past five years the IRS began investing in expert systems technology in order to improve its performance of these mission

---

[13] Department of Commerce; "Artificial Intelligence," *U.S. Industrial Outlook 1989*, pp. 26/3-26/5.

[14] J.S. Congress, Office of Technology Assessment, *Federal Funding for Artificial Intelligence Research and Development*. (For full citation, see section 5 of the bibliography).

tasks.[15] The IRS currently has eighteen expert systems applications in operation or under development.

One of the agency's major applications, the Auditor's Assistant, helps staff classify for audit the citizen and corporate tax returns most apt to result in additional assessments for the IRS. Currently in its fourth year of a five-year project, a real-time pilot implementation of the Auditor's Assistant was field-tested at one service center during the 1990 tax season. The task of classifying tax returns for audit typically is performed as a two-step process. First, tax returns undergo a computerized algorithmic statistical screening. Auditors then manually examine the "sensitive" tax returns to reduce the list further and to identify the key audit issues in the returns that remain. Since audits are extremely time consuming and expensive, auditors attempt, through the manual screening, to identify tax returns with the greatest potential for additional assessments, and that contain the most solid issues to focus on. Experienced auditors who classify the returns use their own set of heuristic criteria to distinguish efficiently between the best and poorest choices.

To develop the Auditor's Assistant, a group of expert auditors from across the country came together every couple of months for several years to participate in knowledge acquisition sessions. In these rigorous sessions, the auditors were not only able to achieve agreement on all primary issues related to tax return classification, but also developed a set of rules that each believed was better than any one individual heuristic. In fact, at the conclusion of the sessions, the auditors insisted that the opportunity to scrutinize their decision-making process with other expert auditors improved their own individual performance. The IRS calculates that use of the Auditor's Assistant will result in an annual $50 million return on investment.

The purpose of another system underway, the Taxpayer's Service Expert Assistant System (TSEAS) is not so much to bring in more revenue as to help the IRS solve a disturbing public relations problem. TSEAS provides guidance to the temporary employees hired by the IRS each tax season to respond to taxpayer phone inquiries. In the past few years, research conducted by the IRS revealed that up to one-third of staff responses to taxpayer phone inquiries were incorrect. Although troubling, the extent of error is not surprising since the information necessary to perform the task is physically dispersed in 159 IRS publications and 10 volumes of tax regulations. After attempts to solve the problem using conventional programming failed, the IRS decided to build an expert system. The design goal is to create an application sufficiently robust to handle about 85 percent of the questions posed by taxpayers.

A pilot implementation that supports accurate, quick, and consistent IRS response to taxpayer phone-in questions was field-tested in Boston during the 1990 tax season. As a result of the expert system, the Boston office's accuracy rate for responses to taxpayers inquiries improved 20 percentage points (from 60 percent to 80 percent) and the overall

---

[15] The IRS investment in expert systems technology began with the selection of a cadre of mid-career staff with strong backgrounds in tax administration to receive training in artificial intelligence. The first group of nine staff were enrolled for two-year terms in AI programs at the University of Pennsylvania, MIT, Columbia, and Stanford. On their return in 1986, the IRS established its AI Laboratory. Today, the Laboratory has expanded to include about two dozen staff.

accuracy rank of the office rose from twenty-third to tenth place (based on 27 national call sites). These figures are expected to improve even more as staff become accustomed to using the system on a full-time basis. No hard data is available yet on how well the system enhanced productivity, but it appears that the application represents a more efficient way to deliver services. IRS plans to expand the field-testing of the system to four additional sites next tax season.

Some of the other IRS applications include: 1) a system, operational since 1988, that processes virtually all of their requests for permission to waive the IRS regulation that requires companies to file a tax return on magnetic media; 2) another system implemented in 1988 that increases the quantity and improves the quality of examiners' letters to taxpayers; 3) a system that determines the actuarial soundness of employee pension plans (the IRS reports that in 1989 this system resulted in a $100 million gain in assessments); 4) a system, scheduled for pilot implementation this year, that helps determine the tax relationship of a subsidiary corporation to a multinational parent company; and, two applications still under development, one that determines whether an individual can be considered an employee for tax purposes, and the other that makes recommendations on the abatement of taxpayer penalties. These and comparable applications demonstrate the IRS' strong confidence in the ability of expert systems technology to perform as a powerful strategic tool in solving modern tax administration problems.

Social Security Administration: The Social Security Administration (SSA) administers the nation's retirement, survivors, and disability insurance programs, as well as its supplemental income program. Approximately 62,000 SSA employees operate out of more than 1,300 district offices and service centers. Claims representatives, SSA's key employees, make eligibility and entitlement determinations in compliance with voluminous regulations that number more than 40,000 pages. These regulations are subject to monthly updates. The claims representatives' training consists of 13-15 weeks of classroom work followed by one year under close supervision in the field. The ability to recall and apply complex logic from the regulations is key to successful performance. Since SSA operates by authority of law, accurate and consistent determinations are vitally important.

SSA created the Expert Systems and Future Technologies Branch (ESFTB) about five years ago to introduce new technologies into agency operations. The ESFTB is actively developing expert decision-support systems for claims representatives. The ultimate goal is to replace the massive paper manuals with an intelligent electronic library of personal-computer-based expert systems fully integrated with the current mainframe support. Progress towards this goal proceeds through incremental development of applications designed to help claims representatives identify the pertinent issues in a claim, provide a step-by-step procedure for arriving at an eligibility determination, and produce both a customized transaction of the criteria used in making the decision and personalized letters to the claimant.

Specific expert systems applications in or awaiting field-testing include: 1) a system designed for use by claims representatives as a training tool and in client interviews for scheduling and executing overpayment waiver claims, 2) a system that produces customized living arrangement notices for supplemental insurance recipients; 3) a system that produces customized continuing review notices for disability beneficiaries; and 4) a system that evaluates gainful activity for disability claims.

In yet another application, SSA is using an outside contractor to develop an expert systems prototype to assist examiners in making disability determinations. The large and complex Disability Program pays out $29 billion in benefits annually. SSA turns over about $750 million annually for States to administer this program. The prototype expert system being developed to modernize this program will be tested and evaluated by disability examiners. SSA expects the prototype to result in more efficient and consistent administration of the Disability Program's laws and regulations.

SSA field personnel estimate that once operational, the expert systems under development will halve the time required to perform specific tasks. They anticipate that expert systems will not only streamline and speed-up case management within the agency, but will also impose a level of consistency on decision-making that will dramatically curtail the claims appeal process, one of the most time-consuming and expensive agency operations.

Office of Management and Budget, Executive Office of the President: The Office of Management and Budget, Executive Office of the President (OMB), administers the federal budget. It also evaluates, formulates, and coordinates management procedures and program objectives in federal departments and agencies. In 1988, OMB created an Expert Systems Unit to improve the management, regulatory, and budget information systems throughout the agency. Manpower reductions and the impending retirement of critical staff prompted the agency to establish an expert systems office to leverage the technological advantages offered by expert systems. This unit identifies suitable applications and disseminates the technology throughout the Office. It also provides technical assistance to users, maintains liaison with industry representatives and other federal agencies, and provides policy and technical assistance on use of the technology to the Executive Office of the President.

OMB administers Executive Branch policy and performs executive oversight with a staff of fewer than five hundred. The Expert Systems Unit currently includes two and one half full-time-equivalent (FTE) staff who receive technical training from a variety of government, academic, and industry sources. The unit's agenda includes developing seven prototype applications intended to provide technical and policy guidance to OMB staff, including: 1) an application on the proper budgetary treatment of federal revenues; 2) an application on the proper budgetary treatment of government transactions involving federal assets; 3) an application for generating custom reports based on parameter settings for computer users; 4) an application for assessing the computer skills of Executive Office staff and formulating recommendations on their training needs; 5) an application to assist support staff in independently updating several federal budget accounting system databases; 6) an application to alert staff to OMB regulations in need of review; and 7) a diagnostic application to aid computer users to solve a constrained set of computer hardware and software problems. OMB considers expert systems technology to be past the research and development stage, and ripe for operational development.

The experience of other federal agencies is consistent with the practices reported in these three cases. For example, the Federal Bureau of Investigation (FBI) has initiated about two dozen projects including: 1) a system that generates personality profiles of violent criminals based on information about a crime and the victim, 2) a system that analyzes incidents of arson in an effort to predict future striking times and locations, and 3) a set of systems designed to assist agents in every type of investigation. The first of the investigative

systems, a counter-terrorism information system, already is operational in several regional offices. The FBI is developing these systems as a way to manage the complexity of modern investigative case management and to reduce the number of years typically needed to train new agents. To accommodate the agents, the investigative system is entirely mouse-driven, requires no typed input, and draws all information from data files. A comparable design will be used for all future investigative systems.[16]

At the Environmental Protection Agency (EPA), staff introduced expert systems technology by conducting an agency-wide survey to identify the most important problems that lent themselves to expert systems solutions. Since the survey, the EPA has fielded or begun developing approximately one dozen applications, primarily to assist in toxic waste management. One application interprets regulatory language governing storage tanks for toxic chemicals. Another uses the technology to calculate the cost and time required to clean up hazardous waste sites. The Expert Disclosure Analysis and Avoidance System (EDAAS), an early application, assists public information specialists to determine which information concerning the manufacture and distribution of toxic chemicals may be released to the public without compromising sensitive data that the EPA must by law keep from public disclosure.

The National Forest Service is producing a number of advisory systems to help foresters make management decisions. For example, the Red Pine Advisor is used in eastern Michigan as a forest administration aid that recommends when to plant and when to harvest. Other systems provide guidance for gypsy moth control or for the use of insecticides with other pests. In future applications, the Forest Service hopes to combine expert systems with geographic information systems to improve fire management.[17]

The growing use of expert systems by the federal government demonstrates the confidence placed in the technology by agencies. Not unlike archives, in recent years the budgets of civilian federal agencies have declined, while the demand for services has risen. These conditions have forced civilian agencies to increase productivity while reducing expenditures. Some federal agencies use the technology primarily to assist in mission tasks, while others concentrate on its benefit with routine administrative problems. The applications involve tasks such as scheduling, planning, selecting, classifying, training, and interpretation - activities performed by archivists in the routine course of their work. Civilian agencies, careful to avoid projects with unproven technologies, have found that modernizing task performance by using expert systems technology with suitable problems represents a key method for managing the "more with less" mandate of the nineties.

---

[16] Information on FBI expert systems is from a presentation made by Gary Gardner (FBI) to the IEEE Technical Committee on Expert Systems Applications (Washington, DC, June 6, 1990) and from Harvey P. Newquist, "Bloodhounds and expert systems," *AI Expert*. (For the full citation, see section 5 of the bibliography).

[17] Norville (1990), p. 22. (For the full citation, see section 5 of the bibliography).

V. The Library Profession's Use of Expert Systems

Many archives are exchanging information on their holdings in online databases that require the use of conventional bibliographic standards. The library profession is exploring expert systems technology chiefly to assist in tasks related to the development of these shared databases, in areas such as cataloging and information retrieval. While a small number of prototypes are being developed for other types of mission functions (e.g., acquisitions or inter-library loan),[18] surprisingly few projects take advantage of the technology for routine administrative tasks, such as planning and scheduling activities, financial analysis, or basic management consultation. Even though operational systems exist in only a handful of libraries, many graduate schools and research libraries are developing prototypes to assess the suitability of the technology to library domains. There are no aggregate estimates on the use of expert systems by the library profession, but their keen interest in the technology is demonstrated through the large number of publications that appear annually in the literature.[19]

Further, a recent three-day conference sponsored by the Graduate School of Library and Information Science at the University of Illinois, "Artificial Intelligence and Expert Systems: Will They Change the Library?", drew more than 100 mid- and upper-level library managers from around the country.[20] Similarly, the Special Libraries Association is devoting a three-day institute to "Intelligent Systems: A Framework for the Future."[21] Moreover, the Council on Library Resources, a private foundation whose grants are often in the forefront of library developments, reports that last year it accepted proposals to grant funds to four expert systems projects while rejecting a number of others.[22] Bibliographic utility vendors also exhibit substantial interest in the technology. OCLC recently subsidized a research project to assess the feasibility of using expert systems to support automatic title-page cataloging. The prototype achieved an 80 percent accuracy rate in performing automatic cataloging from a monograph's title page. Before attempting to build an operational system, the investigators involved in the project recommend further research in four areas, to: 1) achieve a deeper understanding of the task, 2) overcome grammatical irregularities of title page layout, 3) temper scanning and graphical interpretation problems, and 4) minimize problems associated with the complexity of the cataloging rules.[23]

---

[18] Conversation with Professor Linda C. Smith, Graduate School of Library and Information Science, University of Illinois (April 3, 1990).

[19] See *Sources*, section 6.

[20] Conversation with the conference planner, Professor Linda C. Smith: 27th Annual Clinic on Library Applications of Data Processing (March 25-27, 1990), sponsored by the University of Illinois at Urbana-Champaign, Graduate School of Library and Information Science (April 3, 1990).

[21] Held October 22-24, 1990 in Washington, D.C.

[22] Conversation with Eleanor Sacks of the Council on Library Resources (April 4, 1990).

[23] Stuart Weibel, et al., "Automated Title Page Cataloging." (For the full citation, see section 6 of the bibliography).

The three national libraries - the National Library of Medicine, the National Agricultural Library, and the Library of Congress - each approach the technology in a different manner. The Lister Hill Center for Biomedical Communications of the National Library of Medicine (NLM) is the only national library that currently supports a significant research program. MedIndEx, one of NLM's chief research applications, and conceivably the library profession's most advanced expert systems prototype, performs as an indexer's assistant in assigning thesaurus terms to medical literature. The prototype attempts to improve the quality of medical literature indexing by automating to the greatest possible extent index term selection and application of index rules. Progress on the MedIndEx prototype continues with: 1) the development of a more comprehensive knowledge base, 2) augmentation of the user help facility, and 3) additional testing and evaluation of the system.[24]

Apart from the library's research achievements with AI, several years ago NLM's technical services division decided to explore the benefits of using expert systems for core operations. Since summer 1989, the division has devoted a full-time staff position to expert systems development. The goal of the first application is to create a cataloging assistant, with the initial project concentrating on personal name authority control. If successful, this application ultimately will work in conjunction with the division's automated cataloging system to provide expert guidance on determining the authorized form of names. NLM staff selected cataloging as the principal domain for initial expert systems development because it is both extremely labor intensive and entails a long learning-curve to produce experts. Moreover, the staff recognized that knowledge engineering might identify areas where cataloging rules are vague, contradictory, or in need of clarification. In subsequent applications, the staff plans to focus on other aspects of authority work and descriptive cataloging, as well as subject analysis and classification.[25]

The National Agricultural Library (NAL) has focused on operational use of the technology with three systems in use and several others under development. REGIS, an application that includes an embedded expert system to respond to patrons' queries on African aquaculture, is available as a reference aid at NAL and on a floppy disk distributed by the National Technical Information Service (NTIS). A second application provides reference assistance on nutrition directly to patrons. A third serves as an interface that helps users identify and execute a search from a personal computer for information stored in DIALOG (a commercial online bibliographic database service). NAL plans to use expert systems technology for a variety of purposes, including as a front-end navigation tool for multi-media CD-ROMs being compiled on plant information.

Another application underway, CATTUTOR, serves as a tutorial for training inexperienced catalogers. The hypercard application links Machine-Readable Cataloging (MARC) format guidelines and Anglo-American Cataloging Rules, 2nd edition (AACR2) with a menu-driven

---

[24] Susanne Humphrey, Lister Hill National Center for Biomedical Communications (NLM), is the principal developer of MedIndEx.

[25] One other application, MEDSTATS, currently being field tested in the NLM reference division, directs users to statistical sources in the health sciences.

tutorial on descriptive cataloging that instructs and evaluates the performance of trainees. The purpose of the prototype is to shorten the length of time required to develop expert catalogers. Completion of the prototype is scheduled for June 1990. Once completed, the library community will be invited to evaluate the effectiveness of the prototype and to make recommendations for further development.

A final application, the Pesticide Applicator Training (PAT), provides a computer-assisted learning program for individuals who wish to become certified pesticide applicators. This application is being developed by NAL with funding from the Environmental Protection Agency in conjunction with the United States Department of Agriculture (USDA) Extension Service. Every state requires persons who apply pesticides in a commercial setting to attend formal training and/or to pass an exam. PAT will be made available through the training services offered by the agricultural extension service to all these individuals. A recently completed prototype will be demonstrated for the collaborating agencies in 1990.

In 1987, the Library of Congress (LC) established a committee to learn about expert systems technology and to perform a systematic survey within the technical services branch to identify candidate applications. The purpose of the survey was to determine the problem domains most suitable to expert systems technology and to estimate accruable benefits likely to result from the development of operational systems. Technical services tasks were evaluated against a comprehensive list of characteristics identified as "essential" and "highly desirable" for a suitable expert systems domain.[26] Based on this assessment, the committee recommended three application areas from a dozen library domains as the most promising for development: 1) a shelflisting assistant, 2) a series consultant, and 3) a subject-cataloging consultant.

Shelflisting Assistant: Library materials are stored and retrieved through shelflist numbers. Developing a unique alphanumeric representation for an item is largely an algorithmic procedure. Two situations, however, complicate the process: 1) the need to develop a number for an item that has a large number of authors within a particular classification, and 2) the need to develop a number when more than one cataloging item must be taken into account, such as the geographical or subject coverage, in addition to the main-entry heading. In these instances, staff rely on independent judgments formed while referring to classification schedules and related documents to manually determine the shelflist slot. The proposed Shelflisting Assistant is an interactive system that expertly interprets and applies classification schedules and related documentation, and analyzes the patterns and practices implicit in the shelflist itself to create an appropriate call number.

Series Consultant: Most librarians probably consider series work the most problematic aspect of descriptive cataloging. Creating a descriptive entry for a monograph series presents particular problems because: 1) monograph series often are characterized by frequent title changes and numbering peculiarities, 2) the pertinent rules and procedures are numerous and complex, 3) series practices have changed significantly over the years which makes it necessary to relate new items to existing series established with outdated descriptive

---

[26] *Expert Systems: Concepts and Applications*, pp. 23-25. (For the full citation, see section 6 of the bibliography).

practices, and 4) new series descriptions must be integrated into the existing diverse catalog. The Series Consultant would provide guidance to a cataloger in the following broad categories of tasks: 1) establishing new series with proper headings, references, and treatments based on appropriate cataloging rules and procedures, and 2) resolving complex series questions and problems. This application is considered particularly critical because of the scarcity of series cataloging experts.

Subject Cataloging Consultant: When subject cataloging, a librarian makes determinations on the use of terms as headings and references, the precise form of subject headings, the use and form of subdivisions, and the appropriate classification. The performance of this highly skilled task requires that staff have familiarity with a very large body of LC documentation. The proposed Subject Cataloging Consultant would replace all the documentation issued by LC, supporting subject cataloging with an expert system that includes subject heading policies, classification policies, interpretations, and procedures; and that interacts with electronic bibliographic, name and subject authority files, and geographic subdivisions databases. The proposed system would match terms supplied by the cataloger against a thesaurus, display broader and narrower terms, suggest authorized headings, recommend cross-references, and search for appropriate subdivisions. The system also would develop a classification number for each item through use of the primary subject headings or through a thesaurus-assisted search of the classification schedules. The proposers anticipate that use of the Subject Cataloging Consultant would result in greater productivity, improved consistency of practice, and the retention and broad distribution of scarce expertise. LC's three proposals described here appear to be promising applications, but implementation decisions await a detailed cost-benefit analysis.

Although still somewhat speculative, intelligent technologies, in combination with hypermedia, supercomputers, and other new technologies, inspire a rethinking of the conventional library. The computerized card catalog, awaiting the next-generation technology, represents the first step toward the realization of wall-less institutions comprised of computer networks that provide digital access to online information to organizations and individuals around the country. Legislation before the U.S. Congress and strongly endorsed by the Association of Research Libraries (ARL) proposes the development of a National Research and Education Network to support digital database libraries of scientific information that includes in one database climate data, satellite land images, economic materials, and topically-relevant journal articles.[27] The demand for direct electronic access to a variety of discipline-specific information extends beyond the scientific community. Apart from more conventional management functions, expert systems are one tool librarians can use to satisfy patron demands for information packaged and delivered in new ways.

---

[27] Senate Bill 1067: National High Performance Computer Technology Act of 1989. See: Senate Committee on Commerce, Science, and Transportation, *Computer Technology and the New Age of Information*. (For the full citation, see section 7 of the bibliography). Information on ARL endorsement from a phone conversation with Sara Pritchard, ARL Program Officer (April 4, 1990).

VI. State of the Technology

The proficiency achieved in current expert systems technology is the result of nearly thirty years of research. Within the past five years, several key advances mark its maturity into a reliable software tool. First, it is no longer necessary to build an expert system from scratch. A seasoned AI industry has transformed the lessons learned about solving logical problems into structured software shells. In the near term, forecasters predict the appearance of task-oriented shells that will improve the performance of inference engines marketed for specific domains. Second, expert systems, regardless of how they are programmed, typically can now run on many different computers, can be embedded into standard applications, and can process information stored in database management systems. The ability to integrate expert systems technology with existing operations has eliminated one of the key obstacles to widespread adoption.[28] And third, the development of alternate methods of knowledge representation, such as case-based and analogical reasoning, opens the technology to a wider range of domains and more complicated problems.

But probably the most profound breakthrough has just occurred with the recent release by IBM of its new expert systems product, called The Integrated Reasoning Shell (TIRS). According to IBM, TIRS provides new features in addition to functionalities comparable to existing tools, at a fraction of the standard price. Even if TIRS encounters strong market competition, the IBM seal on expert systems technology can be expected to go a long way towards establishing both its legitimacy and standards in the industry.

The advances achieved with expert systems technology, however, tend to be tempered by a fundamental weakness: brittleness. Expert systems only know what they know. They are absolutely useless when asked to solve a problem or reason about a task that falls outside their knowledge base. For instance, a medical diagnosis application may recognize a description of a car covered with rust spots as measles, because its knowledge base does not include the common-sense observation that inanimate objects do not contract infectious diseases. As a result, suitability of the domain to the current technology must be carefully assessed. The most suitable domains rely on narrow, highly structured factual knowledge. For this reason, expert systems are designed to assist, rather than replace staff, freeing experts to concentrate on the more complicated problems within their domain.[29]

Given this limitation, what explains the widespread use of expert systems technology by business, industry, and the government? To answer this question, several years ago Dr. Edward Feigenbaum and his associates at Stanford University, performed a trail-blazing study that examined the effectiveness of 1500 operational business applications. The key finding of their research is the dramatic increase in productivity that results from use of the

---

[28] See for instance, David Freedman, "The big payoff in expert systems"; and "Expert system chip may ease integration," *Computerworld*, for advances in integrating expert systems technology with conventional systems (see section 3 of the bibliography for full citations); see also Kamran Parsaye, et al., *Intelligent Databases: Object-Oriented, Deductive Hypermedia Technologies*, N.Y.: Wiley (c1989).

[29] From keynote address made by Edward Feigenbaum to the Seventh Annual Intelligence Community AI/Advanced Computing Symposium (Reston, Virginia), October 4, 1989.

technology. Productivity was measured in two ways: 1) increases that result from speed-up of task performance and 2) increases that result from transfer of expertise.

Feigenbaum found that assistant-type applications commonly performed professional and semi-professional business tasks at rates twenty, thirty, and even forty times faster than previous operations. Any change that shows an order of magnitude increase in speed of operation (speeds at a factor of 10 greater than the original) is considered revolutionary. In the Feigenbaum study, corporate and industrial managers reported increases in speed from 16 to 400 times greater in tasks using expert systems applications - representing from 1 to 2 orders of magnitude of change. The impact of just one order of magnitude increase in speed is comparable to the difference between driving a car instead of walking, or taking a jet in place of a car to reach a destination.

Further, expert systems technology increased productivity by placing the expertise of experts into the hands of those less skilled performing the same task. Corporate managers contend that for many tasks the top one percent of an organization's staff substantially outperforms the rest of the work force, often exceeding productivity of the bottom half by a factor of twenty.[30] It is expertise - the combination of knowledge, experience, and personal heuristics - that enables the top one percent to arrive at optimum solutions in the least amount of time. Expert systems, built on acquired expertise, transfer the expertise of the top one percent performers to the least productive staff members. Feigenbaum's study found an order of magnitude increase in organizational productivity as a result of expertise transfer, and established that overall operational expert systems generate a fifty-fold increase in productivity for domain tasks.

The striking increases in productivity were not the only benefits reported. Managers assume that the greater consistency imposed on decision-making by expert systems reduces legal vulnerability. Insurance companies, corporate sales departments, and even the IRS and SSA claim that inconsistent treatment of customers (or citizens) encourages legal action. Expert systems also promote organizational stability by capturing and distributing a corporate memory otherwise lost through resignations, retirements, or transfers of critical staff. Further, expert systems expose inconsistent, redundant, or illogical procedures. They are arguably the ultimate tool for learning about a domain and establishing a codified set of rules that represent the logic and heuristics involved in task performance. Further, they offer predictability. The performance of expert systems, unlike humans, does not suffer from burnt-out Fridays.

Moreover, expert systems help organizations manage complexity. Contemporary work situations often require people to master impractical amounts of knowledge. Expert systems, however, can perform favorably with very large volumes of otherwise impenetrable information. For instance, Boeing Corporation decided to develop an expert system to assist mechanics in aircraft repair. Boeing selected this domain because the repair of any one aircraft malfunction typically requires the mechanic's intimate familiarity with at least several feet of manuals. If the amount of information found in several feet of manuals is

---

[30] See for instance, *Augustine's Laws* (pp. 44-45), by Norman Augustine, President and Chief Operating Officer of the Martin Marietta Corporation. (For the full citation, see section 7 of the bibliography).

multiplied by hundreds of possible malfunctions it is apparent why aircraft mechanics need automated intelligent assistance.[31]

Nonetheless, expert systems technology is still an evolving enterprise. During the next few years, forecasters predict that an industry standard for software shells will emerge, especially given the appearance of an IBM tool. This will make it easier for technology consumers to evaluate the tools on the market.[32] Forecasters also anticipate the development of well-defined procedures for managing knowledge engineering projects. Seasoned developers of operational systems appear ready to establish guidelines for iterative development and standard time frames for production. To advance this work the IEEE Technical Committee on Expert Systems Applications sponsored a conference on "Managing Expert Systems Programs and Projects" in the fourth quarter of 1990.

Still needed, however, are verification and validation (v&v) procedures - solid methods for confirming the extent to which a system meets a specification, is internally correct, and solves a real-world problem. The lack of adequate v&v procedures remains a key obstacle limiting the industry's ability to create very large knowledge bases. However, interest by the defense industry and the National Aeronautics and Space Administration (NASA) in autonomous systems - remote systems capable not only of performing a task, but also of monitoring and repairing themselves as necessary - has pushed this issue to the front of the near-term research agenda.

Long-term research addressing the more fundamental problem of application brittleness focuses on two areas: 1) the development of very large "basic" knowledge bases, and 2) machine learning. The industry's key approach to overcoming application fragility is to develop very large "basic" or common sense knowledge bases. The most comprehensive project is a collaboration among leading industry research centers who work together under the rubric of the Microelectronics and Computer Technology Corporation (MCC). MCC has devoted approximately twenty work-years to the development of Cyc, an incremental project to encode the basic rules of physics, cultural belief systems, facts of history, and common activities of humans into a system that will serve as a platform to support specific domain applications. While still in its infancy, Cyc has been used with an application that helps salespersons handling the sale of large computer systems determine customer requirements. Since Cyc contains basic knowledge such as: "rooms in physical structures must be entered and exited through a door," it can assist, for instance, in calculating the amount of cable

---

[31] Remarks by Richard Anderson, Boeing Corporation/Kent Aerospace Center, to representatives from the AI Steering Committee (Seattle: January 25, 1990).

[32] There is quite a bit of activity directed at developing standards for expert systems. The American Institute of Aeronautics and Astronautics (AIAA) is working on a standard terminology glossary, as well as benchmarks for evaluating software tools and tool performance; DARPA is supporting work on the development of knowledge representation standards; and the Institute of Electrical and Electronic Engineers (IEEE) Technical Committee on Expert Systems Applications (ESA) is considering expert systems development standards.

needed to establish a network, given particular information such as room dimensions.[33] In a related effort, Edward Feigenbaum recently announced a national initiative, that includes developers from government, industry, and academia, to codify basic knowledge used in the fields of science and engineering.[34] These platforms, such as Cyc and the science and engineering effort, are considered the chief method for correcting expert system brittleness.

A second long-term research pursuit is machine learning - computer programs that improve with use as a result of actual or simulated experience. Researchers attempt to produce five types of learning with computers: 1) rote learning, where a computer identifies the input of a problem it already has solved and avoids making the same search or calculation mistakes twice, 2) learning from examples, where rules are induced from examples, and the rules from new examples are automatically incorporated into a knowledge base, 3) advice giving/advice taking, where the computer figures out how to break large aggregates of knowledge into processable chunks to integrate them into an existing knowledge base, 4) learning by explanation, where the computer induces rules based on a previous example, case, or theory that solved a problem, and 5) discovery, where computers generate new knowledge from amounts of information too great for human processing, by identifying patterns or anomalies, or by suggesting likely future scenarios based on a stated condition. Discovery constitutes the most sophisticated approach to machine learning. The Central Intelligence Agency (CIA) is using this approach in developing STRADS (Strategic Automated Discovery System), an application that can forecast or "discover" plausible international political, military, and economic events by analyzing an extensive intelligence knowledge base.[35]

New forms of knowledge representation such as case-based reasoning and analogical reasoning expedite the programming of machine learning. For example, a manufacturing application developed by the Lockheed AI Center (Palo Alto, CA.) recommends a method of selection and arrangement for firing and cooling thousands of ceramic pieces of different shapes that minimizes breakage and optimizes production flow. To arrive at an optimal layout, the system considers the size and shape of the various pieces, the furnace load capacity, the intervals at which the pieces are needed for production assembly, and the time required for each load to cool. The expert system first requests information on the materials awaiting firing, then uses case-based reasoning to select from its memory a successful layout that most closely approximates the current reported conditions. The system recommends four different approaches and provides a rationale for each one. The user selects an approach, runs the heating/cooling process, proceeds with production, and then rates the degree to which the approach succeeded or required modification. It is at this point that machine learning occurs, as the system automatically incorporates feedback from the user on the

---

[33] For more information on Cyc, see: Douglas B. Lenat and R.V. Guha. *Building Large Knowledge-Based Systems: Representation and Inference in the Cyc Project.* (For full citation, see section 2 of the bibliography).

[34] From Feigenbaum's presentation at the AISIG '90 research workshop (Washington, DC, Session 8, May 8, 1990).

[35] Dr. Rick Steinhauser (Office of Research and Development, CIA) presented information on STRADS at the AISIG '90 Conference (Washington, DC, Panel 8, May 11, 1990). INLEN, another discovery tool presented at this conference, is being developed by Kenneth Kaufman of George Mason University.

implemented recommendation into its knowledge base. The recommended action with the user's feedback becomes a new case that the system draws on in future dialogue. In this way, the expert system grows more knowledgeable with use.

Modern expert systems technology is sufficiently robust to support a wide range of operational systems, providing five key criteria are met: 1) the application involves a suitable domain, 2) articulate and willing experts are available, 3) developers possess suitable training and experience, 4) end-user involvement is built into development, and 5) executive management supports the enterprise. The absence of any of these criteria will likely shelve an otherwise feasible project.

Suitable domains are characterized by specific task attributes and operational environments. There are four task attributes that should be considered. First, tasks should be sufficiently labor-intensive to assure that an expert system provides a reasonable return on investment. The number of staff and the time it takes to perform a job should warrant automating. Second, the task should require expertise and the need to replicate and distribute it. It probably is not worth the effort to create a system that automates the ordering process at a fast food restaurant. Third, the task should involve judgments that rely on a clearly demarcated set of knowledge and principles. Decisions that hinge primarily on common sense are currently unsuitable for the technology. Fourth, the domain should tolerate "satisficing" solutions - degrees of acceptable answers, rather than only one right answer.

Characteristics of the operational environment should be assessed in conjunction with task attributes. The proposed setting for an expert system application should have at least some of these elements: a shortage of skilled staff to perform a task; an expert's neglect of pressing work due to time spent on more mundane problems; a desire to improve the productivity of staff; a desire to insure consistency of approach and continuity of control between shifts; a need for self-paced educational instruction; a need to preserve or distribute organizational expertise; a need to handle greater levels of complexity in performing tasks; a need to create expertise in an area where not enough knowledge exists; a need to provide expertise in boring or hazardous jobs that do not attract or retain experts; or a need to provide more consistent, timely, reliable, and high-quality decisions.

Once a domain is deemed suitable, the availability of experts should be considered. Most knowledge engineers ask managers to release the staff member(s) they can least afford to do without - the person that everyone performing a task goes to with questions. Moreover, major applications typically require the use of multiple experts. In any expert systems development project, the experts must be given adequate release time to work with the knowledge engineer. Beyond time, it is helpful if domain experts are interested in cognitive processes and the communication of ideas. Since the transfer of the expert's mental model - symbolic networks and patterns of relationships (for instance, metaphors and analogies) they use to understand a problem - to the knowledge engineer is fraught with difficulties, it helps to have a person who likes to figure out the "how" and "why" behind their actions. The expert(s) also should be interested in the project and familiar with its goals from the outset.

Expert systems also require qualified developers. What constitutes a "qualified" developer differs with applications. In general, "personal" stand-alone applications, those that primarily help one person perform one part of a task, have different development requirements than

"organizational" applications, those that support a broader corporate function, are multi-user, and are embedded into or integrated with existing operations. In some instances a domain expert may be the ideal person to build a personal application, but organizational applications invariably require well-trained knowledge engineers. Unfortunately, as a result of the marketing tactics used by many shell vendors, naive users have a tendency to conclude that domain experts, independent of knowledge engineers, are capable of creating all types of expert systems.

Domain experts can be expected to perform effectively as knowledge engineers only in limited circumstances when the following six conditions are met: 1) the proposed application is a "personal" one, requiring only the individual knowledge of the expert, not organizational knowledge drawn from multiple sources in a corporation/agency, 2) the expert is computer literate and a computer enthusiast, 3) the proposed application will be used strictly by the expert, or a small number of staff within the expert's jurisdiction, 4) the method of solving the problem task is characterized by an exceptional degree of expert consensus, 5) the domain lends itself to a low-end stand-alone application that requires less than a few hundred rules, and 6) a less than optimum system is tolerable (users will accept considerable system inefficiency, idiosyncratic logic, and less than peak performance).

For applications that fail to meet these conditions, a qualified knowledge engineer is necessary to reprogram around shell limitations, ascertain the most effective form of knowledge representation and inference for a problem, acquire and model domain knowledge, and integrate the system with existing operations. Knowledge engineers perform highly skilled work, insuring that an application uses the technology and its tools to greatest advantage, serves the appropriate audience, and executes logic in a highly efficient manner.

Finally, the development of a successful expert system application requires end-user involvement as well as executive-level management support. End-user input should inform the design of the application, especially the interface, from project inception. The best projects build end-user investment into the application. No matter what the reason, an application that is rejected by its intended audience constitutes failure. Senior management are an equally important source of support for a project. Their endorsement is vital to secure funds, release necessary staff, and to smooth the roadblocks typically encountered by software development projects.

VII. Implications for Archives

Until recently, human reasoning tasks were considered beyond the scope of automation. But now, expert systems technology offers archival managers a software tool that can be used to solve an entirely new set of organizational problems. The widespread use of the technology by government and industry suggests at least three implications for archivists: 1) a strategy is needed to integrate the technology into the general management of archives, 2) an archival method is needed to appraise existing expert systems, and 3) research is needed to determine an appropriate response to the scholarly use of new "intelligent" research methodologies.

First, the concrete benefits federal agencies, business, and industry enjoy from use of expert systems suggests that archives would enjoy similar benefits by operationalizing the

technology. The studies performed on commercial operational systems clearly demonstrate the technology's broad applicability with both suitable mission-related tasks and routine administrative problems. The increases in productivity alone warrant an evaluation of the suitability of the technology to archival domains. The speed of completing tasks, the distribution of expertise, the increased consistency in decision-making that deters lawsuits or citizen complaints, the customized instructional aids, the retention of corporate memory, and the management of complex functions are all desirable outcomes.

The prospects are excellent that expert systems technology can improve modern archives administration in a wide range of tasks performed in managing records centers to appraising records. To maximize the technology's benefits, archivists need a strategy for developing informed infrastructures, identifying candidate applications, and integrating the technology into existing operations.

Second, the growing use of intelligent technologies by government and institutions raises questions about the appraisal of the structure and content of these systems. Archivists will face the same set of medium-related problems in appraising expert systems as associated with other forms of electronic data. In addition, however, they also will encounter unique problems in evaluating the long-term value of the content of these policy-laden systems. Conventional programs, unlike expert systems, do not regularly incorporate an agency's key policy and interpretive practice into program code. Expert systems knowledge bases by design, however, need a complete set of policy rules with interpretative logic to support task performance. The appearance of "policy-driven" electronic systems promises to pose additional problems for archivists as regulations and procedures that historically appeared in a manager's file begin to appear exclusively as code in an expert system. Indeed, panelists at the 1990 Artificial Intelligence Systems in Government Conference held in May articulated that the tendency for expert systems to serve as the chief receptacle for an agency's policy represented an "obvious trend."[36] As such, archivists need to monitor the use of existing expert systems to determine at what point adjustments in archival practice may be appropriate and necessary.

Finally, the growth in social scientists' use of quantitative research methods in the last several decades suggests that there is a discernable relationship between the emergence of new technologies and scholarly research trends. Social scientists proved quick to adopt quantitative methods of analysis when conventional computing capabilities became available. Currently, researchers are using not only machine-readable numeric records, but have begun to encode textual source materials into electronic form to facilitate new types of computational analysis.

Textual encoding is especially prevalent among humanities scholars. For example, the Data Collection Initiative (DCI) is a far-reaching effort to acquire and make available for research purposes a massive electronic text corpus that will serve as a comprehensive research resource. Sponsored by the Association for Computational Linguistics, the DCI is the most

---

[36] Statements made by panelists at the AI Shells session, AISIG '90 (Washington, DC, May 11, 1990).

extensive international collaboration of its kind. The ultimate goal of the project is to develop a global electronic library available for research use through online terminal access.[37]

Smaller, more focused humanities projects to encode text in machine-readable form are also underway, including: 1) Oxford University's computer database for pre-Restoration English drama, as well as one for texts used by honor students in classics and modern languages, 2) Queens University's electronic library of Canadian literature, and 3) a project by scholars in Toronto and Otago, New Zealand to create a Tudor textbase of written works from the sixteenth century. Humanities scholars predict that the millions of words of text already available in machine-readable form represent only a minute fraction of textual source materials that will be encoded in the next ten to fifteen years.[38]

Concerned about the speed with which researchers are creating textbases for computational research in the absence of agreed upon encoding standards, the Association for Computers and the Humanities, in conjunction with the Association for Computational Linguistics, and the Association for Literary and Linguistic Computing received funding from the National Endowment for the Humanities (NEH) and the European Economic Community to implement a proposal to establish guidelines for the encoding and interchange of machine-readable textual records. The first phase of funding will be devoted to the needs of literary, linguistic, and text-oriented historical research.[39] As with conventional archival records, scholars believe that the reuse of textual databases by scholars other than the original encoders will soon represent - if it does not already - the predominant use. For this reason they have launched a campaign to establish a machine-readable text archive designed to serve a new generation of researchers.[40] NEH awarded Rutgers and Princeton universities a planning grant for this purpose.

The reason textual encoding has significance for archivists is because these efforts suggest the emergence of new scholarly research methods. As massive historical, political, scientific, and literary text bases become available in machine-readable form, "intelligent" text analysis represents the logical progression in research methods.[41] Tomorrow's researchers, computer

---

37 Conversations with Don Walker, Bellcore (May 14, 1990), and Mark Lieberman, AT&T Bell Laboratories (June 5, 1990).

38 *Proposal for Funding for an Initiative to Formulate Guidelines for the Encoding and Interchange of Machine-Readable Text*, by The Association for Computers and the Humanities, the Association for Computational Linguistics, and the Association for Literary and Linguistic Computing, prepared for NEH (1988), p. 12.

39 Ibid., p. 59.

40 Ibid., p. 55.

41 In fact, artificial intelligence is already being used for the purpose of historical analysis. Richard Ennals, in his book *Artificial Intelligence: Applications to Logical Reasoning and Historical Research* (p. 125) reports the study by French social historian Beatrice Henin, on the effect of King Louis XIV's mid-seventeenth century order to expand the city on people living in different quarters of Marseilles. For her research, Henin developed a computer file of leasehold documents drawn up by notaries and property inventories taken at the time of deaths. Towards the end of her research, she became interested in the interior decor

veterans since elementary school, are not likely to sift through thirty feet of paper records when new technologies permit intelligent pattern-matching, anomaly detection, and policy inference with massive volumes of electronic information.[42] In the future, the majority of scholars can be expected to choose computers in combination with intelligent software as the research method of choice. The shift in scholarly research trends toward computer analysis of global libraries of digitized textbases merits the attention of archivists. In fact, it is likely that unless archival practices begin to address the changing research methods of patrons, the prestige, authority, and funding of archives is destined to rapidly diminish.[43]

For three reasons - 1) expert systems promise to improve the management of archives, 2) expert systems pose unique challenges for archival appraisal, and 3) expert systems inaugurate the use of new research methodologies - the archival profession needs a solid working knowledge of the technology, to identify its implications for the management of archives, and to transform archival practice as necessary.

Conclusion

During the next decade, the use of expert systems, especially in combination with other new technologies such as hypermedia and hypertext, supercomputers, digital imaging, and optical disks, will revolutionize the management of archives. The impact of the technology will be felt most strongly by corporate, university, and the spectrum of government archives. For example, Minnesota recently adopted a state-wide strategy for introducing expert systems technology into government operations. The strategy includes a collaborative program with IBM, whereby the state, in conjunction with a local vendor, will develop 48 expert systems that perform basic government functions which IBM will market as customizable shells to the other forty-nine states. The state's governing leadership anticipates that the royalties received from the sale of these systems both will defray initial project costs and support additional expert systems development. To implement the strategy, Minnesota established a ten-person Knowledge Systems Center and instituted new state job classifications and

---

of houses from different social classes. Her use of artificial intelligence to analyze pictures on the walls of rooms, largely with religious themes, led her to develop a new model which can be applied to the understanding of Protestant and Catholic families in England in the 17th century.

[42] In a number of schools in England instructors already are using artificial intelligence to teach students about historical research methods. For instance, AI technologies have been used to develop front-ends that allow middle to high-school age students to query census databases; in the design of computer simulations for events such as the Russian Revolution or the development of the European Economic Community to help students understand the meaning of historical context; or, with machine-readable trade directories from the 19th century to allow twelve year-old students to analyze changes in a community over time. See, Richard Ennals, *Artificial Intelligence: Applications to Logical Reasoning and Historical Research* (pp. 47-63; 75-85, and 96-103).

[43] Avra Michelson (NARA/NSZ) and Jeff Rothenberg (The RAND Corporation) are preparing a paper on the policy implications for archives of the relationship between new technological capabilities and changes in scholarly research trends.

training programs designed to produce the necessary para-technical staff.[44] The Minnesota strategy suggests that state governments are prepared to integrate the technology into existing operations.

Archivists need a deep understanding of expert systems technology before attempting to determine its suitability to archival domains, or its likely affect on core mission functions. The requisite first step is to develop informed infrastructures within archival repositories. Educational programs are needed that encourage managers of existing systems, researchers from AI laboratories, commercial contractors, and allied professionals to share their expertise with archival staff. Once staff members are sufficiently schooled in the technology, they can begin to identify candidate applications for expert systems development. Research prototyping may be desirable for problem domains where feasibility is questionable but the potential payback high. This education, research, and analysis will establish the foundation for developing a strategy to integrate the technology into existing operations and for monitoring the implications of existing systems for archives.

However routine, archives can prepare for expert systems and other new technologies by first, developing an informed infrastructure, second, identifying suitable candidate applications, and third, performing research prototyping with high-return domains. In suitable domains, archivists should collaborate on research and development projects with other archives, government agencies, corporate organizations, or libraries to maximize benefits and defray costs.

The information age is rapidly exploding into the information revolution. Like its nineteenth century predecessor, the information revolution will necessitate radical changes in a number of modern professions. Understanding the implications and planning for the use of new and emerging technologies will equip archivists to deal with an increasingly large segment of their mission.

---

[44] Information on the Minnesota initiative is from a presentation made by Mary Leonard, Manager, Minnesota Knowledge Systems Center, at the panel "State and Local Governments Use of Expert Systems," AISIG '90 (Washington, DC, May 10, 1990).

SOURCES

The chief source used in virtually every section of this report was Edward Feigenbaum, et al., *The Rise of the Expert Company*.

1. Introduction

Details on Fortune 500 expert systems applications (IBM, DuPont, and American Express) is from Edward Feigenbaum, et al., *The Rise of the Expert Company*, and "Whatever happened to AI?", *PC Computing*, 2:3 (March 1989).

Additional information on IBM applications appears in the following: "Putting knowledge to work," *Time* (March 28, 1988), pp. 61-62; and Jeffrey Melaragno and Mary Kay Allen, *A Plan for the Application of Artificial Intelligence to DoD Logistics* (Report PL816R1) Bethesda, Maryland: Logistics Management Institute (October 1989), pp. 1-5, 2-7.

Additional sources used for the section on DuPont include: "Putting knowledge to work." *Time* (March 28, 1988), p. 62; "Keeping pace with perpetual change," *Computerworld* (October 16, 1989), p. 67; and Jeffrey Melaragno and Mary Kay Allen, *A Plan for the Application of Artificial intelligence to DoD Logistics* (Report PL816R1), Bethesda: Maryland: Logistics Management Institute (October 1989), pp. 1/5 & 2/9-210.

Information about American Express is drawn also from: Andrew Kupfer, "Now, live experts on a floppy disk," *Fortune* (October 12, 1987), pp. 69 & 74; and "The big payoff in expert systems," *Computerworld*, Section 2, Focus on Integration (April 3, 1989), p. 44.

For a discussion of the impact of expert systems on mental labor, see Edward Feigenbaum, et al., *The Rise of the Expert Company*, p. 6; and Deborah D. Wolfgram, et al., *Expert Systems for the Technical Professional*, p. 244.

2. What Are Expert Systems?

Every monograph that appears in Section 1 of the bibliography ("Core Texts") explains the distinctions between expert systems and conventional programming. For a fuller explanation of the differences, see especially: Edward Feigenbaum, et al., *The Rise of the Expert Company*; *Building Expert Systems*, ed. by Hayes-Roth, Waterman, and Lenat; Deborah D. Wolfgram et al., *Expert Systems for the Technical Professional*, and David Bendel Hertz, *The Expert Executive*.

For an extensive discussion of suitable domains and examples of operational expert systems see Edward Feigenbaum, et al., *The Rise of the Expert Company*. See also Dimitris N. Chorafas, *Applying Expert Systems in Business*; Paul Harmon and David King, *Expert Systems: Artificial Intelligence in Business*; Anna Hart, *Expert Systems: An Introduction for Managers*; Deborah D. Wolfgram et al., *Expert Systems for the Technical Professional*; Nancy R. Miller, *Supercomputers and Artificial Intelligence: Federal Initiatives*; "Now, live experts on a floppy disk," *Fortune* (October 12, 1987); "Putting knowledge to work," *TIME* (March 28, 1988); and Daniel Todd, "Expert systems go retail," *InformationWEEK* (November 6, 1989).

3. History of Expert Systems

The sources referred to in writing this section include: Edward Feigenbaum et al., *The Rise of the Expert Company*; Edward Feigenbaum and Pamela McCorduck, *The Fifth Generation*; Anna Hart, *Expert Systems: An Introduction for Managers*; Deborah Wolfgram et al., *Expert Systems for the Technical Professional*; and Library of Congress, Congressional Research Service, *The Impact of Information Technology on Science*.

4. Building Expert Systems

The key source for this section was Feigenbaum, et al., *The Rise of the Expert Company*. For additional information on building expert systems, knowledge engineering, programming languages, hardware, shells, and development costs, see: Paul Harmon and David King, *Expert Systems: Artificial Intelligence in Business*. David W. Rolston's *Principles of Artificial Intelligence and Expert Systems Development* is another good source.

For additional information on software shells, see: Paul Harmon's monthly newsletter, *Expert System Strategies*, and his *Expert Systems Tools and Applications*; Deborah D. Wolfgram, et al., *Expert Systems for the Technical Professional*; and Anna Hart, *Expert Systems: An Introduction for Managers*.

For further discussion of AI hardware, see: Deborah D. Wolfgram, *Expert Systems for the Technical Professional*; on types and sizes of expert systems, see: Paul Harmon, et al., *Expert Systems Tools and Applications*, and Stephen Ruth and Christopher K. Carlson, "Low end expert systems in business: high yields for stand-alone applications," *Office: Technology and People* (1989).

On development costs, see: Guy Benchimol, et al., *Developing Expert Systems*; on problems acquiring expert knowledge, see: Dianne C. Berry, "The problem of implicit knowledge," *Expert Systems* (August 1987), as well as the citations that appear in section 4 of the bibliography ("Knowledge Acquisition").

5. Federal Government Use of Expert Systems

The military use of expert systems is documented in Nancy R. Miller, *Supercomputers and Artificial Intelligence: Federal Initiatives*, p. 3; and Deborah D. Wolfgram et al., *Expert Systems for the Technical Professional*, p. 258.

The section on civilian federal expert systems was compiled from on-site and phone interviews with staff in federal agencies, and from secondary sources. In federal agencies for which an expert systems contact was not known, NARA staff began gathering information by contacting the science and technology offices listed in the *U.S. Government Manual*. Approximately twenty federal agencies were contacted. In all of these agencies, staff reported at least some use of expert systems technology, and that their agencies considered expert systems a mainstream technology like any other software tool.

Details on IRS expert systems applications is taken from: "Symposium: Expert Systems." Susan G. Hadden, ed., *Journal of Policy Analysis and Management*; Valerie Norville, "Expert Systems," *Federal Computer Week* (July 24, 1989); Richard K. Schreiber, "Artificial Intelligence Applications in the IRS," *The Inference Engine*; "Bailing out taxpayer assistance," *Computerworld*; and Cynthia Whipple, et al., "Knowledge acquisition for an Internal Revenue Service classification system." This section also refers to a presentation made by Rick Schreiber to the IEEE Task Force on Expert Systems (WDC: July 25, 1989) and conversations with IRS AI staff Rick Schreiber (September 22, 1989 & April 24, 1990), Jim Needham (September 28, 1989), and Bruce Ramsey (October 10, 1989).

Information on Social Security Administration applications is taken from "Symposium: Expert Systems," Susan G. Hadden, ed. *Journal of Policy Analysis and Management*; and from conversations with Bruce Whaite, ES&FT, on March 20, 1990, and Sue Peters, ES&FT, on March 23, 1990.

The description of OMB applications was based on discussions with Eric Won, Manager, Expert Systems Unit, OMB, on October 13, 1989 and March 27, 1990. Additional information as well as a prototype demonstration was provided by Arlene Dell, Budget Management and Information Specialist on March 27, 1990. Eric Won also provided NARA with copies of unpublished materials on the OMB expert systems program that were used in this section.

The two key sources for information on EPA expert systems applications are Valerie Norville's articles entitled "Expert systems", *Federal Computer Week* (July 24, 1989 & February 26, 1990); for a description of EDAAS, see: Jerald L. Feinstein and Frederick Siems, "EDAAS: an expert system at the US Environmental Protection Agency for avoiding disclosure of confidential business information," *Expert Systems*; and Jerald L. Feinstein, "A knowledge-based expert system used to prevent the disclosure of sensitive information at the United States Environmental Protection Agency," in *Computing Power and Legal Reasoning*.

6. Library Profession's Use of Expert Systems

Publications on the library profession's use of expert systems are prolific. See for instance: *Expert Systems in Reference Service*; *Expert Systems in Libraries*; *Information Processing and Management*, 23/2 (1987): Special issue on Expert Systems and Library and Information Science; Gilbert K. Krulee and Alexander Vrenios, "An expert system model of a reference librarian," *Library Software Review*; P.F. Anderson, "Expert systems, expertise, and the library and information professions," *Library and Information Science Research*; Roy Meador III and Glenn R. Wittig, "Expert systems for automatic cataloging based on AACR2: a survey of research," *Information Technology and Libraries*; Donald T. Hawkins, "Applications of artificial intelligence (AI) and expert systems for online searching," *Online*; H.M. Brooks, "Expert systems and intelligent information retrieval," *Information Processing and Management*; Harold Borko, "Getting started in library expert systems research," *Information Processing and Management*; and Alina Vickery and Helen Brooks, "Expert systems and their applications in LIS," *Online Review*.

The sources referred to for NLM's MedIndEex are Susanne M. Humphrey, "A knowledge-based expert system for computer-assisted indexing," *IEEE Expert*; and S.M. Humphrey and N.E. Miller, "Knowledge-based indexing of the medical literature: the Indexing Aid Project," *Journal of the American Society for Information Science*. Sally Sinn, Deputy Chief of Technical Services Division, and Paul Weiss, Expert Systems Staff (Technical Services) provided information on the operational applications at NLM on April 2, 1990.

The details on expert systems applications at NAL are from phone conversations with Gary McCone, Database Administration Branch, (NAL) on April 3, 1990; Claudia Weston, CATTUTOR developer, on April 4, 1990; and Sarah Thomas, Head of Technical Services, on April 17, 1990. The NAL staff disagree whether CATTUTOR and the Pesticide Applicator Training application should be considered expert systems. NARA decided to include them in this report for two reasons: 1) because PAT was developed with an expert systems shell, and 2) because the use of expert knowledge and reasoning rather than algorithms is core to both applications. For information on the CD-ROM expert system front-end, see: Pamela R. Mason, "Planning a multimedia CD-ROM," *National Online Meeting: Proceedings 1990*.

The source for LC's survey of candidate applications is from *Expert Systems: Concepts and Applications*, prepared by Charles Fenly, Library of Congress in association with Howard Harris, RMG Consultants, Inc. Additional information was provided by Charles Fenly in a phone conversation on April 5, 1990.

7. State of the Technology

The near-term appearance of task-oriented shells is widely anticipated. Steven W. Oxman (principal at the OXKO Corporation, Annapolis, Maryland) presented this forecast at his workshop on knowledge acquisition at the annual meeting of the International Association of Knowledge Engineers (College Park, Md.) on June 26, 1989, as did Dr. John Boose at his presentation to representatives from the AI Steering Committee at Boeing's Advanced Technology Center (Seattle) on January 25, 1990.

The discussion of new forms of knowledge representation, including case-based reasoning and reasoning by analogy, is from notes taken on the AI Steering Committee Tour of West Coast AI research laboratories (January 22-25, 1990). Information on TIRS is from Lance B. Eliot, "TIRS: the great blue hope." *AI Expert*.

Findings from the Feigenbaum study referred to throughout this section are taken from *The Rise of the Expert Company* or from notes taken at Feigenbaum's keynote address to the Seventh Annual Intelligence Community AI/Advanced Computing Symposium (Reston, Virginia), October 4, 1989.

The section on Cyc is from notes taken at a briefing with Ramanathan Guha, MCC (Stanford University), January 22, 1990. See also: Douglas B. Lenat and R. V. Guha, *Building Large Knowledge-Based Systems: Representation and Inference in the Cyc Project*; and David Coursey, "Scientist finds common sense for computing in gossip tabloid," *MIS Week*. The discussion of machine-learning is from notes of conversation with David Rolston, ESL (Sunnyvale, CA) on January 23, 1990.

The section that includes criteria for developing successful applications refers to Jeffrey Melaragno and Mary Kay Allen, *A Plan for the Application of Artificial Intelligence to DoD Logistics*; David Rolston, *Principles of Artificial Intelligence and Expert Systems Development*; and Paul Harmon and David King, *Expert Systems*.

BIBLIOGRAPHY

Note: If you can read only one book on expert systems, read Edward Feigenbaum, et al., *The Rise of the Expert Company.*

1. Core Texts

Building Expert Systems, (Teknowledge series in knowledge engineering; v. 1), edited by Frederick Hayes-Roth, Donald A. Waterman, and Douglas B. Lenat, Reading, Massachusetts: Addison-Wesley Publishing Company, Inc. (1983).

Feigenbaum, Edward A., and Pamela McCorduck, *The Fifth Generation: Artificial Intelligence and Japan's Computer Challenge to the World*, Reading, Massachusetts: Addison-Wesley Publishing Company (1983).

Feigenbaum, Edward, Pamela McCorduck, and H. Penny Nii, *The Rise of the Expert Company*, New York: TimesBooks (1988).

Harmon, Paul, and David King, *Expert Systems: Artificial Intelligence in Business*, New York: John Wiley & Sons, Inc. (1985).

Harmon, Paul, Rex Maus, and William Morrissey, *Expert Systems Tools and Applications*, New York: John Wiley & Sons (1988).

Rolston, David W, *Principles of Artificial Intelligence and Expert Systems Development*, New York: McGraw-Hill Book Company (1988).

Wolfgram, Deborah D., Teresa J. Dear, and Craig S. Galbraith, *Expert Systems for the Technical Professional*, New York: John Wiley & Sons (1987).

2. General Works: Monographs

Aleksander, Igor, and Piers Burnett, *Thinking Machines: The Search for Artificial Intelligence*, New York: Alfred A. Knopf (1987).

Applied Expert Systems, E. Turban and P.R. Watkins, eds. North-Holland: Elsevier Science Publishers B.V. (1988).

Benchimol, Guy, Pierre Levine, and Jean-Charles Pomerol, *Developing Expert Systems*, London: Instrument Society of America (1987).

Bowerman, Robert G., and David E. Glover, *Putting Expert Systems into Practice*, New York: Van Nostrand Reinhold Company (1988).

Chorafas, Dimitris N., *Applying Expert Systems in Business*, New York: McGraw-Hill Book Company (1987).

Drechsler, Frank, and John Bateson, *Management, Information and Expert Systems*, Dublin: Irish Academic Press (1986).

Ennals, Richard, *Artificial Intelligence: Applications to Logical Reasoning and Historical Research*, Chichester: Ellis Horwood (1985).

Expert Systems for Business, edited by Barry G. Silverman, Reading, Massachusetts: Addison-Wesley Publishing Company (1987).

Forsyth, Richard, and Chris Naylor, *The Hitchhiker's Guide to Artificial Intelligence*, London: Chapman and Hall/Methuen (1986).

Frenzel, Louis E., Jr., *Understanding Expert Systems*, Indianapolis, Indiana: Howard W. Sams & Company (1987).

Gardin, J. C., *Artificial Intelligence and Expert Systems: Case Studies in the Knowledge Domain of Archaeology*, Chichester: Ellis Horwood (1988).

Hart, Anna, *Expert Systems: An Introduction for Managers*, Columbia, Maryland: GP Publishing, Inc. (1988).

Hertz, David Bendel, *The Expert Executive: Using AI and Expert Systems for Financial Management, Marketing, Production, and Strategy*, New York: John Wiley & Sons (1988).

Lenat, Douglas B., and R.V. Guha; *Building Large Knowledge-Based Systems: Representation and Inference in the Cyc Project*, Reading, Massachusetts: Addison-Wesley Publishing Company (c1990).

Managing Artificial Intelligence and Expert Systems, Edited by Daniel A. De Salvo and Jay Liebowitz. Englewood Cliffs, New Jersey: Prentice-Hall (c1990).

Rothenberg, Jeff, Jody Paul, Iris Kameny, James R. Kipps, and Marcy Swenson; *Evaluating Expert System Tools: A Framework and Methodology*, Prepared for the Defense Advanced Research Projects Agency, Santa Monica, CA: Rand Corporation (1987).

Siegel, Paul, *Expert Systems: A Non-Programmers Guide to Development and Applications*, Blue Ridge Summit, Pa.: Tab Books, Inc. (1986).

Waterman, Donald A., *A Guide to Expert Systems*, Reading, Massachusetts: Addison-Wesley Publishing Company (1986).

Winograd, Terry, and Fernando Flores; *Understanding Computers and Cognition: A New Foundation for Design*, Norwood, New Jersey: Ablex Publishing Corporation (1986).

3. General Works: Periodical Articles

Badalamenti, Anthony F., "Sigmund Freud on artificial intelligence (an artificial interview)", *Applied Artificial Intelligence* (Jan.-Mar. 1988) 2/1, pp. 65-89.

Burgess, John, "The electronic experts," *The Washington Post*, Business section (October 1, 1989) pp. H-1, H-4.

Coursey, David, "Scientist finds common sense for computing in gossip tabloid," *MIS Week* (March 19, 1990), p. 37.

Eliot, Lance B., "Information management and expert systems: a pedagogical approach", *Information Management Review* 2:3 (Winter 1987), pp. 63-69.

_____, "TIRS: the great blue hope," *AI Expert* (April 1990), pp. 9-11.

_____, "The underlying beat of the industry," *AI Expert*, (November 1989), pp. 9-11.

"Expert system chip may ease integration," *Computerworld*, 23:44 (October 30, 1989), pp. 41-45.

"Expert systems go retail," *InformationWEEK*, (November 6, 1989), p. 28.

Fordyce, Kenneth, Peter Norden, and Gerald Sullivan; "Review of expert systems for the management science practitioner," *Interfaces* 17:2 (March-April 1987), pp. 64-77.

Freedman, David, "The big payoff in expert systems," *Computerworld*: Focus on Integration, Section 2, (April 3, 1989), pp. 36-44.

Gibson, Stanley, "Reality checks in at AI show," *Computerworld*: (September 4, 1989), pp. 23 & 26.

Goul, Michael, "On building expert systems for strategic planners: a knowledge engineer's experience," *Information & Management 12* (1987), pp. 131-141.

Grimm, Vanessa Jo, "Congress keeps eye on computing in investigations," *Government Computer News* 8:12 (June 12, 1989), pp. 146.

Harris, Larry R., "Getting started with expert systems," *Management Information Systems Week* (June 5, 1989), p. 31.

Hawkins, Donald T., "Artificial intelligence (AI) and expert systems for information professionals -- basic terminology," *Online* 11:5 (September 1987), pp. 91-98.

"Keeping pace with perpetual change," *Computerworld* (October 16, 1989), pp. 61, 66-67.

Keyes, Jessica, "Why expert systems fail," *AI Expert* (November 1989), pp. 50-53.

Leary, Michael, and Agustin Rodriguez-Bachiller; "Expertise, domain-structure and expert system design: a case study in development control," *Expert Systems* 6:1 (February 1989), pp. 18-23.

Muir, Bonnie M., "Trust between humans and machines, and the design of decision aids," *International Journal of Man-Machine Studies* 27 (1987), pp. 527-539.

Newquist, Harvey P. III, "LISP chips to chocolate chips," *AI Expert* (November 1989), pp. 55-58.

"A maturing AI is finding its way in the world," *Computerworld* (February 19, 1990), p. 23.

"Next-generation race bogs down," *IEEE Spectrum*, (June 1987), pp. 28-33.

"Now, live experts on a floppy disk," *Fortune*, (October 12, 1987), pp. 69-82.

"Picking the right expert system application," *AI Expert* (September 1989), pp. 44-47.

"Putting knowledge to work," *Time* (March 28, 1988), pp. 60-63.

Ruth, Stephen R. and Christopher K. Carlson; "Low-end expert systems in business: high yields for stand-alone applications," *Office: Technology and People*, 4:4 (1989), pp. 299-312.

Shipley, Chris, "Whatever happened to AI?" *PC Computing* 2:3 (March 1989), pp. 64-74.

"Smart factories: America's turn?" *Business Week*, (May 8, 1989), pp. 142-145, 148.

Smith, Karl A., "Educational engineering: heuristics for improving learning effectiveness and efficiency," *Engineering Education* (February 1987), pp. 274-279.

Travis, Irene L., "Knowledge representation in artificial intelligence," Unpublished paper, c1990.

Turkle, Sherry, "Artificial intelligence and psychoanalysis: a new alliance," *Daedelus* 117 (Winter 1988), pp. 241-268.

Waltz, David L., "The prospects for building truly intelligent machines," *Daedelus* 117 (Winter 1988), pp. 191-212.

"Whatever happened to AI?" *PC Computing* 2:3 (March 1989), pp. 64-74.

Woods, William A., "Important issues in knowledge representation," *Proceedings of the IEEE* 74:10 (October 1986), pp. 1322-1334.

4. Knowledge Acquisition

Alexander, James H., et. al, "Ontological analysis: an ongoing experiment," *International Journal of Man-Machine Studies*, 26 (1987), pp. 473-485.

Belkin, N.J., and H.M. Brooks; "Knowledge elicitation using discourse analysis," *International Journal of Man-Machine Studies*, 27 (1987), pp. 127-144.

Berry, Dianne C., "The problem of implicit knowledge," *Expert Systems*, 4/3 (August 1987), pp. 144-151.

Bylander, Tom, and B. Chandrasekaren; "Generic tasks for knowledge-based reasoning: the "right" level of abstraction for knowledge acquisition," *International Journal of Man-Machine Studies*, 26 (1987), pp. 231-243.

Cleaves, David A., "Cognitive biases and corrective techniques: proposals for improving elicitation procedures for knowledge-based systems," *International Journal of Man-Machine Studies*, 27 (1987), pp. 155-166.

Cooke, Nancy M., and James E. McDonald; "A formal methodology for acquiring and representing expert knowledge," *Proceedings of the IEEE*, 74:10 (October 1986), pp. 1422-1430.

Cullen, J. and Bryman, A.; "The knowledge acquisition bottleneck: time for reassessment?" *Expert Systems*, 5:3 (August 1988), pp. 216-225.

Davies, Miranda, and Simon Hakiel; "Knowledge harvesting: a practical guide to interviewing," *Expert Systems*, 5:1 (February 1988), pp. 42-50.

Eliot, Lance B., "Investigating the nature of expertise: analogical thinking, expert systems, and ThinkBack," *Expert Systems*, 4:3 (August 1987), pp. 190-195.

Enkawa, Takao, and Gavriel Salvendy; "Underlying dimensions of human problem solving and learning: implications for personnel selection, training, task design and expert system," *International Journal of Man-Machine Studies*, 30 (1989), pp. 235-254.

Gaines, Brian R., "An overview of knowledge acquisition and transfer," *International Journal of Man-Machine Studies*, 26 (1987), pp. 453-472.

Garg-Janardan, Chaya, and Gavriel Salvendy; "A conceptual framework for knowledge elicitation," *International Journal of Man-Machine Studies*, 26 (1987), pp. 521-531.

_____, "A structured knowledge elicitation methodology for building expert systems," *International Journal of Man-Machine Studies*, 29 (1988), pp. 377-406.

Gruber, Thomas R., "Acquiring strategic knowledge from experts," *International Journal of Man-Machine Studies*, 29/5 (1988), pp. 579-597.

Haywood, S. A., "Structured analysis of knowledge," *International Journal of Man-Machine Studies*, 26 (1987), pp. 487-498.

Johnson, Paul E., "Specification of expertise," *International Journal of Man-Machine Studies*, 26 (1987), pp. 161-181.

Kitto, Catherine M., and John Boose, "Heuristics for expertise transfer: an implementation of a dialog manager for knowledge acquisition," *International Journal of Man-Machine Studies*, 26 (1987), pp. 183-202.

Kornell, Jim, "Formal thought and narrative thought in knowledge acquisition," *International Journal of Man-Machine Studies*, 26 (1987), pp. 203-212.

LaFrance, Marianne, "The knowledge acquisition grid: a method for training knowledge engineers," *International Journal of Man-Machine Studies*, 26 (1987), pp. 245-255.

Littman, David C., "Modelling human expertise in knowledge engineering: some preliminary observations," *International Journal of Man-Machine Studies*, 26 (1987), pp. 81-92.

Olson, Judith Reitman, and Henry H. Rueter, "Extracting expertise from experts: methods for knowledge acquisition," *Expert Systems*, 4:3 (August 1987), pp. 152-168.

Regoczei, Stephen, "Creating the domain of discourse: ontology and inventory," *International Journal of Man-Machine Studies*, 27 (1987), pp. 235-250.

Woods, David D., "Mapping cognitive demands in complex problem-solving worlds," *International Journal of Man-Machine Studies*, 26 (1987), pp. 257-275.

5. Expert Systems in Government

"Bailing out taxpayer assistance," *Computerworld* (September 11, 1989), pp. 39 & 49.

Feinstein, Jerald L., "A knowledge-based expert system used to prevent the disclosure of sensitive information at the United States Environmental Protection Agency," *Computing Power and Legal Reasoning*, edited by Charles Walter. St. Paul, Minnesota: West Publishing Company (1985), pp. 681-697.

Feinstein, Jerald L., and Frederick Siems, "EDAAS: an expert system at the U.S. Environmental Protection Agency for avoiding disclosure of confidential business information," *Expert Systems*, 2:2 (April 1985), pp. 72-85.

Franklin, Jude, and Kamal Karna; "Expert systems in government," *IEEE Expert* (Spring 1989), pp. 3-5.

BEST COPY AVAILABLE

Melaragno, Jeffrey, and Mary Kay Allen; *A Plan for the Application of Artificial Intelligence to DoD Logistics* (Report PL816R1) Bethesda, Maryland: Logistics Management Institute (October 1989).

Newquist, Harvey P., "Bloodhounds and expert systems," *AI Expert* (March 1990), pp. 67-70.

Norville, Valerie, "Expert systems," *Federal Computer Week* (July 24, 1989) pp. 30, 34-36; and article by same title (February 26, 1990), pp.22-24, 32.

Robb, David W., "Expert-system shell helps public find USDA data," *Government Computer News* (October 2, 1989), p. 15.

Schreiber, Richard K., "Artificial intelligence applications in the IRS," *The Inference Engine*, 1:2 (September 1989), pp. 4-6.

"Symposium: Expert Systems," Susan G. Hadden, editor; *Journal of Policy Analysis and Management*, 8:2 (1989), pp. 182-208; includes: Jerald L. Feinstein, "Introduction to expert systems"; W. James Hadden, Jr., "Expert systems for environmental permits"; Gerald J. Grady and Richard K. Schreiber, "Expert systems applications in the IRS"; Ed Leary, "Expert systems at the Social Security Administration"; and Susan G. Hadden, "The future of expert systems in government."

Taft, Darryl K., "Weather service's next mainframe may be its last," *Government Computer News* (April 2, 1990), p. 25.

Whipple, Cynthia, Lawrence Davis, John Kam, and Jim Needham; "Knowledge acquisition for an Internal Revenue Service classification system," Unpublished paper (August 1988).

U.S. Congress, Office of Technology Assessment, Communication and Information Technologies Program; *Federal Funding for Artificial Intelligence Research and Development*, Staff Paper (June 1987).

6. Library and Archival Applications

Aluri, Rao, and Donald E. Riggs; "Application of Expert Systems to Libraries," *Advances in Library Automation and Networking*, 2 (1988), pp. 1-43.

Anderson, P.F., "Expert systems, expertise, and the library and information professions," *Library and Information Science Research*, 10:4 (Oct/Dec 1988), pp. 367-368.

Borko, Harold, "Getting started in library expert systems research," *Information Processing & Management*, 23:2 (1987), pp. 81-87.

Brooks, H. M., "Expert systems and intelligent information retrieval," *Information Processing & Management*, 23:4 (1987), pp. 367-382.

Croft, W. Bruce, "Approaches to intelligent information retrieval," *Information Processing & Management*, 23:4 (1987), pp. 249-254.

Curtin, Bonnie Rose, "Is conservation ready for artificial intelligence?" *The Abbey Newsletter*, 14:1 (February 1990), pp. 1-2.

De Salvo, Daniel A., Amy E. Glamm, and Jay Liebowitz; "Structured design of an expert system prototype at the National Archives," Chapter 3 in *Expert Systems for Business*, edited by Barry G. Silverman. Reading, Massachusetts: Addison-Wesley Publishing Company (1987), pp. 40-78.

De Salvo, Daniel A., and Jay Liebowitz; "The application of an expert system for information retrieval at the National Archives," *Telematics and Informatics*, 3/1 (1986), pp. 25-38.

"Expert Systems: Concepts and Applications," Prepared by Charles Fenly, Library of Congress in association with Howard Harris, RMG Consultants, Inc., *Advances in Library Information Technology*, Issue Number 1, Washington, D.C.: Cataloging Distribution Service, Library of Congress (1988).

"Expert Systems and Library and Information Science," special issue of *Information Processing and Management*, 23:2 (1987).
Expert Systems in Libraries, edited by Rao Aluri and Donald E. Riggs; Norwood, New Jersey: Ablex (1990).

Expert Systems in Reference Service, Christine Roysdon and Howard D. White, editors. New York: Haworth Press, (c1989).

45 BEST COPY AVAILABLE

Hawkins, Donald T., "Applications of artificial intelligence (AI) and expert systems for online searching," *Online* (January 1988), pp. 31-43.

Hirtle, Peter B., "Artificial intelligence, expert systems, and archival automation," *Provenance*, 5 (1987), pp. 76-88.

Humphrey, S.M., and N.E. Miller; "Knowledge-based indexing of the medical literature: the Indexing Aid Project," *Journal of the American Society for Information Science*, 38:3 (May 1987), pp. 184-196.

Humphrey, Susanne M., "A knowledge-based expert system for computer-assisted indexing," *IEEE Expert* (Fall 1989), pp. 25-38.

Jaussaud, Renee, "The archivist's assistant: from the expert's perspective," Chapter 4 in *Managing Artificial Intelligence and Expert Systems*, Daniel A. De Salvo and Jay Liebowitz, eds. Englewood Cliffs, New Jersey: Yourdon Press (c1990), pp. 53-79.

Krulee, Gilbert K. and Alexander Vrenios; "An expert system model of a reference librarian," *Library Software Review*, 8 (Jan/Feb 1989), pp. 13-16.

Mason, Pamela R., "Planning a multimedia CD ROM," *National Online Meeting: Proceedings 1990*, (New York, May 1-3, 1990), edited by Martha E. Williams, Medford, N.J.: Learned Information, Inc. (1990), pp. 265-271.

Meador, Roy, and Glenn R. Wittig; "Expert systems for automatic cataloging based on AACR2: A survey of research," *Information Technology and Libraries* (June 1988), pp. 166-171.

O'Neill, Margaret, and Anne Morris; *The Role of Information Scientists in the Development of Expert Systems: Interim Report*, British Library Research Paper 57, West Yorkshire, Great Britain: British Library Board (1989).

Rau, Lisa F., "Knowledge organization and access in a conceptual information system," *Information Processing & Management*, 23:4 (1987), pp. 269-283.

Shoval, Peretz, "Principles, procedures and rules in an expert system for information retrieval," *Information Processing & Management*, 21:6 (1985) pp. 475-487.

Vedder, Richard G., Maurice G. Fortin, Scott A. Lemmermann, and Ralph N. Johnson; "Five PC-Based Expert Systems for Business Reference: An Evaluation," *Information Technology and Libraries*, (March 1989) pp. 42-54.

Vickery, Alina, and Helen Brooks; "Expert systems and their applications in LIS," *Online Review*, 11:3 (1987), pp. 149-165.

Waters, Samuel T., "Answerman, the expert information specialist: An expert system for retrieval of information from library reference books." *Information Technology and Libraries*, (September 1986), pp. 204-212.

Weibel, Stuart, Michael Oskins, and Diane Vizine-Goetz; "Automated title page cataloging: a feasibility study," Unpublished paper, OCLC Office of Research (February 19, 1988), 27p. Forthcoming in *Information Processing and Management*.

7. Other Sources

Augustine, Norman R., *Augustine's Laws*, New York: Penguin Books (1987).

Department of Commerce, Technology Administration, *Emerging Technologies: A Survey of Technical and Economic Opportunities*, (Spring 1990), 55pp.

Department of Commerce, *U.S. Industrial Outlook 1989*, pp. 26-3 - 26-5.

Guinan, Patricia J., "Patterns of Excellence for IS Professionals: An Analysis of Communication Behavior," Washington, D.C., *International Center for Information Technologies Press*, (1988).

The Impact of the Information Age on Science, Hearings before the Task Force on Science Policy of the Committee on Science and Technology, U.S. House of Representatives, Ninety-ninth Congress, first session, September 10-12, 1985, Washington: U.S. G.P.O., 1986 (Y 4.Sci 2:99/103).

The Impact of the Information Age on Science, Report prepared by the Congressional Research Service, Library of Congress; Transmitted to the Task Force on Science Policy, Committee on Science and Technology, U.S. House of Repreeentatives, Ninety-ninth Congress, second session, Washington: U.S. G.P.O., 1986 (Y 4.Sci 2:99/103).

Miller, Nancy R., "Supercomputers and artificial intelligence: federal initiatives," Library of Congress, Congressional Research Service, Science Policy Research Division, Issue Brief, Order code IB85105; Archived 5/19/87, Updated 3/22/87.

National Research Council, *The National Challenge in Computer Science and Technology*, Washington: National Academy Press (1988).

Perolle, J. A., "Artificial intelligence in the workplace," *The Worker in Transition: Technological Change*, edited by T.J. Kozik and D.G. Jansson, Presented at: The Worker in Transition: Technological Change Conference (Bethesda: Md., April 4-7, 1989). c.1989, pp. 35-41.

_____, "Intellectual assembly lines: the rationalization of managerial, professional, and technical work," *Computers and the Social Sciences*, 2:3 (July-September 1986), pp. 111-121.

_____, "The social impact of computing: ideological themes and research issues," *Social Science Computer Review*, 6:4 (Winter 1988), pp. 469-480.

Senate Committee on Commerce, Science, and Transportation, Subcommittee on Science, Technology, and Space, Computer Technology and the New Age of Information, Hearing (September 15, 1989).

Spradley, James P., *The Ethnographic Interview*, New York: Holt, Rinehart and Winston (1979).

"Strategic computing: a status report," *IEEE Spectrum*, (April 1987), pp. 50-54.

"Symposium on the impact of computers on science and society," *IRE Transactions on Electronic Computers*, Volume EC-5/3 (September 1956), pp. 142-158.

*U.S. Government Printing Office: 1991 — 291-631